arXiv:1612.00547v2 [math.OC] 27 Jan 2017

Gradient Descent Efficiently Finds the Cubic-Regularized Non-Convex Newton Step

Yair Carmon John C. Duchi
Stanford University
{yairc ,jduchi}@stanford.edu

Abstract

We consider the minimization of non-convex quadratic forms regularized by a cubic term, which exhibit multiple saddle points and poor local minima. Nonetheless, we prove that, under mild assumptions, gradient descent approximates the *global minimum* to within ε accuracy in $O(\varepsilon^{-1}\log(1/\varepsilon))$ steps for large ε and $O(\log(1/\varepsilon))$ steps for small ε (compared to a condition number we define), with at most logarithmic dependence on the problem dimension. When we use gradient descent to approximate the Nesterov-Polyak cubic-regularized Newton step, our result implies a rate of convergence to second-order stationary points of general smooth non-convex functions.

1 Introduction

We study the optimization problem

$$\underset{x\in\mathbb{R}^d}{\text{minimize}} \;\; f(x) \triangleq \frac{1}{2}x^\top Ax + b^\top x + \frac{\rho}{3}\|x\|^3, \tag{1}$$

where the matrix A is symmetric and possibly indefinite. The problem (1) arises in Newton's method with cubic regularization, first proposed by Nesterov and Polyak [18]. The method consists of the iterative procedure

$$x_{t+1} = \underset{x\in\mathbb{R}^d}{\operatorname{argmin}} \left\{ \nabla g(x_t)^\top (x - x_t) + \frac{1}{2}(x - x_t)^\top \nabla^2 g(x_t)(x - x_t) + \frac{\rho}{3}\|x - x_t\|^3 \right\} \tag{2}$$

for (approximately) minimizing a general smooth function g, requiring sequential solutions of problems of the form (1). The Nesterov-Polyak scheme (2) falls into the broader framework of trust-region methods [6, 4], which methods require repeated solution of problems similar to (1), though they often permit approximate solutions and adaptively choose the regularization parameter ρ. Such methods are among the most practically successful and theoretically sound approaches to non-convex optimization [6, 18, 4]. Indeed, Nesterov and Polyak establish convergence rates for convergence to second-order stationary points of g.

Standard methods for solving the problem (1) exactly require either factorization or inversion of the matrix A. However, the cost of these operations scales poorly with the problem dimensionality. In contrast, matrix-free methods, which access A only through matrix-vector products, often scale well to high dimensions and leverage structure in A (c.f. [23]), particularly when A is a Hessian. Even without special structure, the product $\nabla^2 g(x)v$ often admits the finite difference approximation $\delta^{-1}(\nabla g(x + \delta v) - \nabla g(x))$, which requires only two gradient evaluations. In neural networks and

other arithmetic circuits, back-propogation-like methods allow exact computation of Hessian-vector products at a similar cost [20, 21]. It is thus of interest to explore matrix-free methods guaranteed to solve (1) efficiently.

1.1 Outline of our contribution

In this paper we consider gradient descent, which is perhaps the simplest matrix-free optimization method. We show that for the problem (1), under mild assumptions, gradient descent converges to the *global minimum* in spite of the multiple saddle points and suboptimal local minima of the objective f. Moreover, we provide convergence rates close to the established rates of gradient descent for smooth convex optimization.

We begin our development in Section 2 with a number of definitions and results, specifying our assumptions, characterizing the solution to problem (1), and proving that gradient descent converges to the global minimum of f. Additionally, we show that gradient descent produces iterates with monotonically increasing norm. This property is essential to our results, and we use it extensively throughout the paper.

In Section 3.1 we provide non-asymptotic rates of convergence for gradient descent, which are our main results. We outline our proofs in Section 4, deferring technical arguments to appendices as necessary. Our first convergence guarantee includes the term $\log(1/|v_1^T b|)$, where v_1 is the eigenvector corresponding to the smallest eigenvalue of A. When this quantity is infinite—the so-called "hard case" for non-convex quadratic problems [6]—we show that gradient descent can solve a slightly perturbed problem (with high probability), exhibiting convergence rates scaling no worse than logarithmically in problem dimension and achieving accurate solutions for the true problem. Our results have close connections with the convergence rates of gradient descent on smooth convex functions and of the power method, which we discuss in Section 7.

We illustrate our results with a number of experiments, which we report in Section 3.2. We explore the trajectory of gradient descent on non-convex problem instances, demonstrating its dependence on problem conditioning and the presence of saddle points. We then illustrate our convergence rate guarantees by running gradient descent over an ensemble of random problem instances. This experiment suggests the sharpness of our theoretical guarantees.

In Section 5 we extend our scope to step sizes chosen by exact line search. If the search is unconstrained, the method may fail to converge to the global minimum, but success is guaranteed for a guarded variation of exact line search. Unfortunately, we have thus far been unable to give rates of convergence for this scheme, though its empirical behavior is at least as strong as standard gradient descent.

As our primary motivation for solving problem (1) is the Nesterov-Polyak method (2), in Section 6 we consider a method for minimizing a general non-convex function g, which approximates the iterations (2) via gradient descent. When g has 2ρ-Lipschitz continuous Hessian, we show that this method finds a point y such that $\|\nabla g(y)\| \leq \epsilon$ after roughly ϵ^{-2} gradient evaluations, which is the rate for gradient descent applied directly on g [17, Ex. 1.2.3]. Additionally, our procedure provides the second-order guarantee $\nabla^2 g(y) \succeq -\sqrt{\rho\epsilon}I$, and thus we give a first-order method with non-asymptotic convergence guarantees to second-order stationary points at essentially no additional cost over gradient descent.

1.2 Related work

A number of researchers have considered low-complexity methods methods for solving the problem (1). Cartis et al. [4] propose solution methods working in small Krylov subspaces, and Bian-

concini et al. [2] apply a matrix-free gradient method (NMGRAD) in conjunction with an early stopping criterion for the problem. Both approaches exhibit strong practical performance, and they enjoy first-order convergence guarantees for the overall optimization method (in which problem (1) is an iteratively solved sub-problem). In both works, however, it appears challenging to give convergence guarantees for the iterative subproblem solvers, and they do not provide the second-order convergence rates of Nesterov and Polyak's Newton method (2). In their paper [5], Cartis et al. give sufficient conditions for a low-complexity approximate subproblem solution to guarantee such second-order rates, but it is unclear how to construct a first-order method fulfilling the conditions.

There is a large literature on the quadratic trust region problem [6, 9, 10, 7], where one replaces the regularizer $(\rho/3)\|x\|^3$ with the constraint $\|x\| \leq R$; we give only a necessarily limited overview. Classical low complexity methods for this problem include subspace methods and the Steihaug-Toint truncated conjugate gradient method; we know of no convergence guarantees in either case. Tao and An [22] give an analysis of projected gradient descent with a restart scheme that guarantees convergence to the global minimum; however, the number of restarts may be proportional to problem dimension, suggesting potential difficulties for large-scale problems.

More recently, Hazan and Koren [11] give a first-order method that solves the trust-region problem with an accelerated rate. They find an ε-suboptimal point for the trust region problem in $O(1/\sqrt{\varepsilon})$ matrix-vector multiplies (ignoring logarithmic factors and problem-dependent constants) by reducing the trust-region problem to a sequence of approximate eigenvector problems. Ho-Nguyen and Kılınç-Karzan [12] provide a different perspective, showing how a single eigenvector calculation can be used to reformulate the non-convex quadratic trust region problem into a convex QCQP, efficiently solvable with first-order methods.

Concurrent to our work, Agarwal et al. [1] show an accelerated rate of convergence for the cubic problem (1) via reductions to fast approximate matrix inversion and eigenvector computations. Their rates of convergence are better than those we achieve when ε is large relative to problem conditioning. However, while these works indicate that solving (1) is never harder than approximating the bottom eigenvector of A, the regime of linear convergence we identify shows that it is sometimes much easier. In addition, we believe that our results provide interesting insights on the potential of gradient descent and other direct methods—those without explicit eigenvector calculations or matrix inversions—for non-convex problems.

Another related line of work is the study of the behavior of gradient descent around saddle-points and its ability to escape them [8, 14, 15]. A common theme in these works is an "exponential growth" mechanism that pushes the gradient descent iterates away from critical points with negative curvature. This mechanism plays a prominent role in our analysis as well, highlighting the implications of negative curvature for the dynamics of gradient descent.

2 Preliminaries and basic convergence guarantees

We begin by defining some (mostly standard) notation. Our problem (1) is to solve

$$\underset{x\in\mathbb{R}^d}{\text{minimize}} \;\; f(x) \triangleq \frac{1}{2}x^\top A x + b^\top x + \frac{\rho}{3}\|x\|^3,$$

where $\rho > 0$, $b \in \mathbb{R}^d$ and $A \in \mathbb{R}^{d\times d}$ is a symmetric (possibly indefinite) matrix, and $\|\cdot\|$ denotes the Euclidean norm. The eigenvalues of the matrix A are $\lambda^{(1)}(A) \leq \lambda^{(2)}(A) \leq \cdots \leq \lambda^{(d)}(A)$, where any of the $\lambda^{(i)}(A)$ may be negative. We define the eigengap of A by $\mathsf{gap} \triangleq \lambda^{(k)}(A) - \lambda^{(1)}(A)$ where k is the first eigenvalue of A strictly larger than $\lambda^{(1)}(A)$. The eigendecomposition of A is

$A = \sum_{i=1}^{d} \lambda^{(i)}(A) v_i v_i^T$ where the vectors $v_i \in \mathbb{R}^d$ are orthonormal. For any vector $w \in \mathbb{R}^d$ we let

$$w^{(i)} = v_i^T w \text{ denote the } i\text{th coordinate of } w \text{ in the eigenbasis of } A.$$

We let $\|\cdot\|_{\text{op}}$ be the ℓ_2-operator norm, so $\|A\|_{\text{op}} = \max_{u:\|u\|=1} \|Au\|$, and define

$$\gamma \triangleq -\lambda^{(1)}(A), \quad \gamma_+ \triangleq \gamma \vee 0, \quad \text{and} \quad \beta \triangleq \|A\|_{\text{op}} = \max\{|\lambda^{(1)}(A)|, |\lambda^{(d)}(A)|\},$$

so that the function f is non-convex if and only if $\gamma > 0$ (and is convex when $\gamma_+ = 0$). We remark that our results continue to hold when β is an upper bound on $\|A\|_{\text{op}}$ rather than its exact value. We say that a function g is *L-smooth* on a convex set X if $\|\nabla g(x) - \nabla g(y)\| \le L \|x - y\|$ for all $x, y \in X$; this is equivalent to $\left\|\nabla^2 g(x)\right\|_{\text{op}} \le L$ for Lebesgue almost every $x \in X$ and is equivalent to the bound $|g(x) - g(y) - \nabla g(y)^T (x - y)| \le \frac{L}{2} \|x - y\|^2$ for $x, y \in X$.

2.1 Characterization of f and its global minimizers

Throughout the paper, we let s denote a solution to problem (1), *i.e.* a global minimizer of f, and define the matrix

$$A_s \triangleq A + \rho \|s\| I,$$

where I is the $d \times d$ identity matrix. The minimizer s has the characterization [18, Section 5] as the solution to the equality and inequality

$$\nabla f(s) = A_s s + b = 0 \quad \text{and} \quad \rho \|s\| \ge \gamma, \tag{3}$$

and s is unique whenever $\rho \|s\| > \gamma$. We may write the gradient and Hessian of f as

$$\nabla f(x) = Ax + b + \rho \|x\| x = A_s(x - s) - \rho(\|s\| - \|x\|)x \text{ and } \nabla^2 f(x) = A + \rho \|x\| I + \rho \frac{xx^T}{\|x\|}.$$

The optimal value of f admits the expression and bound

$$f(s) = \frac{1}{2} s^T A s + b^T s + \frac{\rho \|s\|^3}{3} = -\frac{1}{2} s^T A_s s - \frac{\rho \|s\|^3}{6} \le -\frac{\rho \|s\|^3}{6}, \tag{4a}$$

and, using the fact that $s^T A_s s = -b^T s \le \|b\| \|s\|$, we derive the lower bound

$$f(s) \ge -\frac{1}{2} \|b\| \|s\| - \frac{\rho \|s\|^3}{6}. \tag{4b}$$

Algebraic manipulation also show that

$$f(x) = f(s) + \frac{1}{2} (x - s) A_s (x - s)^\top + \frac{\rho}{6} (\|s\| - \|x\|)^2 (\|s\| + 2\|x\|), \tag{5}$$

which makes it clear that s is indeed the global minimum, as both of the x-dependent terms are positive and minimized at $x = s$, and the minimum is unique whenever $\|s\| > \gamma/\rho$, because $A_s \succ 0$ in this case.

The global minimizer admits the following equivalent characterization whenever the vector b is not orthogonal to the eigenspace associated with $\lambda^{(1)}(A)$.

Claim 2.1. *If $b^{(1)} \neq 0$, s is the unique solution to $\nabla f(s) = 0$ and $b^{(1)} s^{(1)} \le 0$.*

Proof. Let s' satisfy $\nabla f(s') = 0$ and $b^{(1)}s'^{(1)} \le 0$. Focusing on the first coordinate of the first condition, we have $0 = [\nabla f(s')]^{(1)} = (-\gamma + \rho\|s'\|)s'^{(1)} + b^{(1)}$. Therefore, $b^{(1)} \neq 0$ implies both $s'^{(1)} \neq 0$ and $-\gamma + \rho\|s'\| \neq 0$. This strengthens the inequality $b^{(1)}s'^{(1)} \le 0$ to $b^{(1)}s'^{(1)} < 0$. Hence $-\gamma + \rho\|s'\| = -b^{(1)}s'^{(1)}/[s'^{(1)}]^2 > 0$ and, by the characterization (3) of s, if a critical point satisfies $\rho\,\|s'\| > \gamma$ then it is the unique global minimum s. □

The norm of s plays an important role in our analysis, so we provide a number of bounds on it. First, observe that whenever $\|s\| > \gamma/\rho$, we have $\|s\| = \|A_s^{-1}b\| \le (-\gamma + \rho\|s\|)^{-1}\|b\|$. Solving for $\|s\|$ thus gives the upper bound

$$\|s\| \le \frac{\gamma}{2\rho} + \sqrt{\left(\frac{\gamma}{2\rho}\right)^2 + \frac{\|b\|}{\rho}} \le \frac{\beta}{2\rho} + \sqrt{\left(\frac{\beta}{2\rho}\right)^2 + \frac{\|b\|}{\rho}} \triangleq R \tag{6a}$$

where we recall that $\beta = \|A\|_{\mathrm{op}} \ge |\gamma|$. An analogous lower bound on $\|s\|$ is available: we have $\|s\| \ge \gamma/\rho$, and if $b^{(1)} \neq 0$, then $\|s\| = \|A_s^{-1}b\| \ge |b^{(1)}|/(-\gamma + \rho\|s\|)$ implies

$$\|s\| \ge \frac{\gamma}{2\rho} + \sqrt{\left(\frac{\gamma}{2\rho}\right)^2 + \frac{|b^{(1)}|}{\rho}} \ge -\frac{\beta}{2\rho} + \sqrt{\left(\frac{\beta}{2\rho}\right)^2 + \frac{|b^{(1)}|}{\rho}}. \tag{6b}$$

We can also prove a different lower bound with the similar form

$$\|s\| \ge R_c \triangleq \frac{-b^TAb}{2\rho\|b\|^2} + \sqrt{\left(\frac{b^TAb}{2\rho\|b\|^2}\right)^2 + \frac{\|b\|}{\rho}} \ge -\frac{\beta}{2\rho} + \sqrt{\left(\frac{\beta}{2\rho}\right)^2 + \frac{\|b\|}{\rho}}. \tag{7}$$

The quantity R_c is the *Cauchy radius* [6]—the magnitude of the (global) minimizer of f in the subspace spanned by b: $R_c = \mathrm{argmin}_{\zeta\in\mathbb{R}} f(-\zeta b/\|b\|)$. To see the claimed lower bound (7), set $x_c = -R_c b/\|b\|$ and note that $f(x_c) = -(1/2)\|b\|R_c - (\rho/6)R_c^3$. Therefore, $0 \le f(x_c) - f(s) \le \frac{1}{2}\|b\|(\|s\| - R_c) + \frac{1}{6}\rho(\|s\|^3 - R_c^3)$, which implies $\|s\| \ge R_c$.

2.2 Properties and convergence of gradient descent

The gradient descent method begins at some initialization $x_0 \in \mathbb{R}^d$ and generates iterates via

$$x_{t+1} = x_t - \eta\nabla f(x_t) = (I - \eta A - \rho\eta\,\|x_t\|\,I)x_t - \eta b, \tag{8}$$

where η is a fixed step size. Recalling the definitions (6a) and (7) of R and R_c as well as $\|A\|_{\mathrm{op}} = \beta$, throughout our analysis we make the following assumptions.

Assumption A. *The step size η in* (8) *satisfies* $0 < \eta \le \frac{1}{4(\beta+\rho R)}$.

Assumption B. *The initialization of* (8) *satisfies* $x_0 = -r\frac{b}{\|b\|}$, *with* $0 \le r \le R_c$.

We begin our treatment of the convergence of gradient descent by establishing that $\|x_t\|$ is monotonic and bounded (see Appendix A for a proof).

Lemma 2.2. *Let Assumptions A and B hold. Then the iterates* (8) *of gradient descent satisfy* $x_t^T\nabla f(x_t) \le 0$, *the norms* $\|x_t\|$ *are non-decreasing, and* $\|x_t\| \le R$.

This lemma is the key to our analysis throughout the paper. The next lemma shows that x_t and b have opposite signs at all coordinates in the eigenbasis of A.

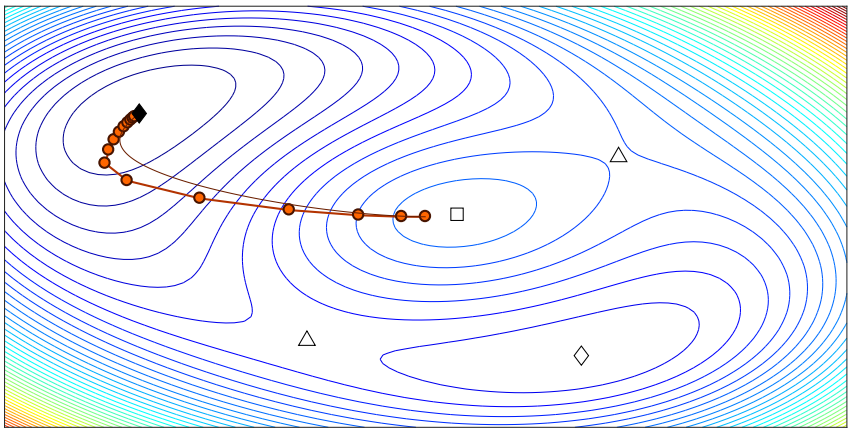

Figure 1. Contour plot of a two-dimensional instance of (1), with $A = [-7, 1; 1, -7]$, $b = [0.2; 0]$ and $\rho = 1$. The brown line is the path of the continuous time process $\dot{x} = -\nabla f(x)$; the line of circles indicates the path of gradient descent with $\eta = 0.35$. Note the presence of local maxima ($\square$), saddle points ($\triangle$), and local minima ($\Diamond$).

Lemma 2.3. *Let Assumptions A and B hold. For all $t \geq 0$ and $i \in \{1, ..., d\}$*

$$x_t^{(i)} b^{(i)} \leq 0, \;\; b^{(i)} s^{(i)} \leq 0, \;\; \text{and } x_t^{(i)} s^{(i)} \geq 0.$$

Consequently, $x_t^T b \leq 0$ and $x_t^T s \geq 0$ for every t, and $s^T b \leq 0$.

Proof. We first show that $x_t^{(i)} b^{(i)} \leq 0$. Writing the gradient descent recursion in the eigen-basis of A, we have

$$x_t^{(i)} = \left(1 - \eta \lambda^{(i)}(A) - \eta\rho\|x_{t-1}\|\right) x_{t-1}^{(i)} - \eta b^{(i)}. \tag{9}$$

Assumption A and Lemma 2.2 imply $1 - \eta\lambda^{(i)}(A) - \eta\rho\|x_{t-1}\| \geq 1 - \eta(\beta + \rho R) > 0$ for all t, i. Therefore, $x_t^{(i)} b^{(i)} \leq 0$ if $x_0^{(i)} b^{(i)} \leq 0$; the initialization in Assumption B guarantees this. To show $b^{(i)} s^{(i)} \leq 0$, we use the fact that $b = -A_s s$ to write

$$b^{(i)} s^{(i)} = -\left(\lambda^{(i)}(A) + \rho\|s\|\right)[s^{(i)}]^2 \leq 0$$

as $\lambda^{(i)}(A) + \rho\|s\| \geq 0$ for every i by the condition (3) defining s.

Next, we see that $b^{(i)} s^{(i)} = -(\lambda^{(i)}(A) + \rho\|s\|)[s^{(i)}]^2 \leq 0$, as $b = -A_s s$ and $\lambda^{(i)}(A) + \rho\|s\| \geq 0$ for by the condition (3) defining s. Multiplying $x_t^{(i)} b^{(i)} \leq 0$ and $b^{(i)} s^{(i)} \leq 0$ yields $x_t^{(i)} s^{(i)} [b^{(i)}]^2 \geq 0$. The coordinate-wise update (9) and Assumption B show that $b^{(i)} = 0$ implies $x_t^{(i)} = 0$ for every t, and therefore $x_t^{(i)} s^{(i)} \geq 0$. $\square$

Lemmas 2.2, 2.3, and Claim 2.1 immediately lead to the following guarantee.

Proposition 2.4. *Let Assumptions A and B hold, and assume that $b^{(1)} \neq 0$. Then $x_t \to s$ and $f(x_t) \downarrow f(s)$ as $t \to \infty$.*

Proof. By Lemma 2.2, the iterates satisfy $\|x_t\| \leq R$ for all t. Since $\left\|\nabla^2 f(x)\right\|_{\text{op}} \leq \beta + 2\rho\|x\|$, the function f is $\beta + 2\rho R$-smooth on the set $\{x \in \mathbb{R}^d : \|x\| \leq R\}$ containing all the iterates x_t.

Therefore, by the definition of smoothness and the gradient step,

$$f(x_{t+1}) \le f(x_t) - \eta\|\nabla f(x_t)\|^2 + \frac{\eta^2}{2}(\beta + 2\rho R)\|\nabla f(x_t)\|^2 \le f(x_t) - \frac{\eta}{2}\|\nabla f(x_t)\|^2 \ ,$$

where final inequality used Assumption A that $\eta \le \frac{1}{4(\beta+\rho R)}$. Consequently, $f(x_t)$ is decreasing and for every $t > 0$,

$$\frac{\eta}{2}\sum_{\tau=0}^{t-1}\|\nabla f(x_\tau)\|^2 \le f(x_0) - f(x_t) \le f(x_0) - f(s).$$

Let s' be any limit point of the sequence x_t (there must be at least one, as the sequence x_t is bounded). The above display implies $\lim_{t\to\infty}\nabla f(x_t) = 0$ and therefore $\nabla f(s') = 0$ by continuity of ∇f. By Lemma 2.3, $x_t^{(1)} b^{(1)} \le 0$ for every t and therefore, $s'^{(1)} b^{(1)} \le 0$. Claim 2.1 thus implies that s' is the unique global minimizer s. We conclude that s is the only limit point of the sequence x_t, whence $x_t \to s$ by boundedness and $f(x_t) \to f(s)$ by continuity of f. □

To handle the case $b^{(1)} = 0$, let $k \ge 1$ be the first index for which $b^{(k)} \ne 0$ (if no such k exists then $b = 0$ and $x_t = 0$ for all t). Let $\hat{s}$ be the unique global minimum of f in the subspace spanned by the eigenvectors of A with eigenvalues $\lambda^{(k)}(A), ..., \lambda^{(d)}(A)$. In this case [4, Sec. 6.1], we have $\|\hat{s}\| \le \|s\|$, since $\hat{s} \ne s$ only if $\rho\|\hat{s}\| \le \gamma$. Applying Lemma 2.2 and Proposition 2.4 in this subspace, we have $\|x_t\| \uparrow \|\hat{s}\|$, and we obtain

Corollary 2.5. *Let Assumptions A and B hold. For all $t \ge 0$, the iterates* (8) *of gradient descent satisfy $x_t^T \nabla f(x_t) \le 0$, the norms $\|x_t\|$ are non-decreasing and satisfy $\|x_t\| \le \|s\|$, and f is $\beta + 2\rho\|s\|$-smooth on a ball containing the iterates x_t.*

Figure 1 provides a graphical representation of these results, showing gradient descent's iterates on an instance of problem (1) exhibiting numerous stationary points.

3 Non-asymptotic convergence rates

Proposition 2.4 shows the convergence of gradient for the cubic-regularized (non-convex) quadratic problem (1). We now present stronger non-asymptotic guarantees, including a randomized scheme solving (1) in all cases. We follow with simulations explicating our theoretical results.

3.1 Theoretical results

Our primary result, Theorem 3.1, gives a convergence rate for gradient descent in the case that $b^{(1)} \ne 0$. In the statements of our results, recall that s is the global minimizer of $f(x) = x^T Ax/2 + b^T x + \rho\|x\|^3/3$, that $\gamma = -\lambda^{(1)}(A)$, $\beta = \|A\|_{\text{op}}$, gap is the eigengap of A, and that $\gamma_+ = \gamma \vee 0$.

Theorem 3.1. *Let Assumptions A and B hold, $b^{(1)} \ne 0$, and $\varepsilon > 0$. Then $f(x_t) \le f(s) + \varepsilon$ for all*

$$t \ge T_\varepsilon \triangleq \frac{1}{\eta}\left(\tau_{\text{grow}}(b^{(1)}) + \tau_{\text{converge}}(\varepsilon)\right)\min\left\{\frac{1}{\rho\|s\| - \gamma}, \frac{10\|s\|^2}{\varepsilon}\right\} \tag{10}$$

where

$$\tau_{\text{grow}}(b^{(1)}) = 6\log\left(1 + \frac{\gamma_+^2}{4\rho|b^{(1)}|}\right) \text{ and } \tau_{\text{converge}}(\varepsilon) = 6\log\left(\frac{(\beta + 2\rho\|s\|)\|s\|^2}{\varepsilon}\right).$$

In addition, if $\rho\|s\| - \gamma \le \frac{\varepsilon}{10\|s\|^2}$, *then* $f(x_t) \le f(s) + \varepsilon$ *for all*

$$t \ge T'_\varepsilon \triangleq \frac{1}{\eta}\left(\tau_{\text{grow}}(b^{(1)}) + \tau_{\text{converge}}(\varepsilon)\right)\sqrt{\frac{10\|s\|^2}{\varepsilon} \cdot \frac{1}{\mathsf{gap} \wedge \rho\|s\|}}.$$

See Section 4.1 for a proof.

Theorem 3.1 shows that the rate of convergence changes from roughly $O(1/\varepsilon)$ to $O(\log(1/\varepsilon))$ as ε decreases, with an intermediate gap-dependent rate of $O(1/\sqrt{\varepsilon})$. The terms τ_{grow} and τ_{converge} correspond to a period (τ_{grow}) in which $\|x_t\|$ grows exponentially until reaching the basin of attraction to the global minimum and a period (τ_{converge}) of linear convergence to s. Exponential growth occurs only in non-convex problem instances, as $\tau_{\text{grow}} = 0$ when the problem is convex.

The dependence of our result on $|b^{(1)}|$ is unavoidable: if $b^{(1)} = 0$, then gradient descent always remains in a subspace orthogonal to the first (bottom) eigenvector of A, while $s^{(1)}$ might be non-zero; this is the "hard case" of non-convex quadratic problems [6]. We use a small random perturbation to guarantee $|b^{(1)}| \neq 0$ except with negligible probability, which yields the following high probability guarantee, whose proof we provide in Section 4.2.

Theorem 3.2. *Let Assumptions A and B hold,* $\varepsilon, \delta > 0$, *and let* q *be uniformly distributed on the unit sphere in* $\mathbb{R}^d$. *Let* $\tilde{x}_t$ *be generated by the gradient descent iteration* (8) *with* $\tilde{b} = b + \sigma q$ *replacing* b, *where*

$$\sigma = \frac{\rho\varepsilon}{\beta + 2\rho\|s\|} \cdot \frac{\overline{\sigma}}{12} \text{ with } \overline{\sigma} \le 1.$$

Then with probability at least $1 - \delta$, *we have* $f(\tilde{x}_t) \le f(s) + (1 + \overline{\sigma})\varepsilon$ *for all*

$$t \ge \tilde{T}_\varepsilon \triangleq \frac{1 + \overline{\sigma}}{\eta}\big(\tilde{\tau}_{\text{grow}}(d, \delta, \overline{\sigma}) + \tilde{\tau}_{\text{converge}}(\varepsilon)\big)\min\left\{\frac{1}{\rho\|s\| - \gamma}, \frac{10\|s\|^2}{\varepsilon}\right\}, \tag{11}$$

where

$$\tilde{\tau}_{\text{grow}}(d, \delta, \overline{\sigma}) \triangleq 6\log\left(1 + \mathbb{I}_{\{\gamma > 0\}}\frac{3\sqrt{d}}{\overline{\sigma}\delta}\right) \text{ and } \tilde{\tau}_{\text{converge}}(\varepsilon) \triangleq 14\log\left(\frac{(\beta + 2\rho\|s\|)\|s\|^2}{\varepsilon}\right).$$

In addition, if $\rho\|s\| - \gamma \le (1 - 2\overline{\sigma}/3)\frac{\varepsilon}{10\|s\|^2}$, *then* $f(\tilde{x}_t) \le f(s) + (1 + \overline{\sigma})\varepsilon$ *for all*

$$t \ge \tilde{T}'_\varepsilon \triangleq \frac{1 + \overline{\sigma}}{\eta}\left(\tilde{\tau}_{\text{grow}}(d, \delta, \overline{\sigma}) + \tilde{\tau}_{\text{converge}}(\varepsilon)\right)\sqrt{\frac{10\|s\|^2}{\varepsilon} \cdot \frac{1}{\mathsf{gap} \wedge \rho\|s\|}}.$$

To facilitate later discussion, we define $L_s \triangleq \beta + 2\rho\|s\|$; then f is L_s-smooth on the Euclidean ball of radius $\|s\|$. Recalling the definition (6a) of R and noting that $\beta + \rho R \le 2\beta + \rho\|s\| \le 2L_s$ (see also the bound (6b)), we have the following corollary.

Corollary 3.3. *Let the conditions of Theorem 3.2 hold,* $\eta = \frac{1}{4(\beta + \rho R)}$ *and* $\overline{\sigma} = 1$. *Then with probability at least* $1 - \delta$, *we have* $f(\tilde{x}_t) \le f(s) + \varepsilon$ *for all*

$$t \ge \tilde{T}_\varepsilon = O(1) \cdot \min\left\{\frac{L_s}{\rho\|s\| - \gamma}, \frac{L_s\|s\|^2}{\varepsilon}\right\}\log\left[\left(1 + \mathbb{I}_{\{\gamma > 0\}}\frac{d}{\delta}\right)\frac{L_s\|s\|^2}{\varepsilon}\right].$$

We conclude the presentation of our main results with a few brief remarks.

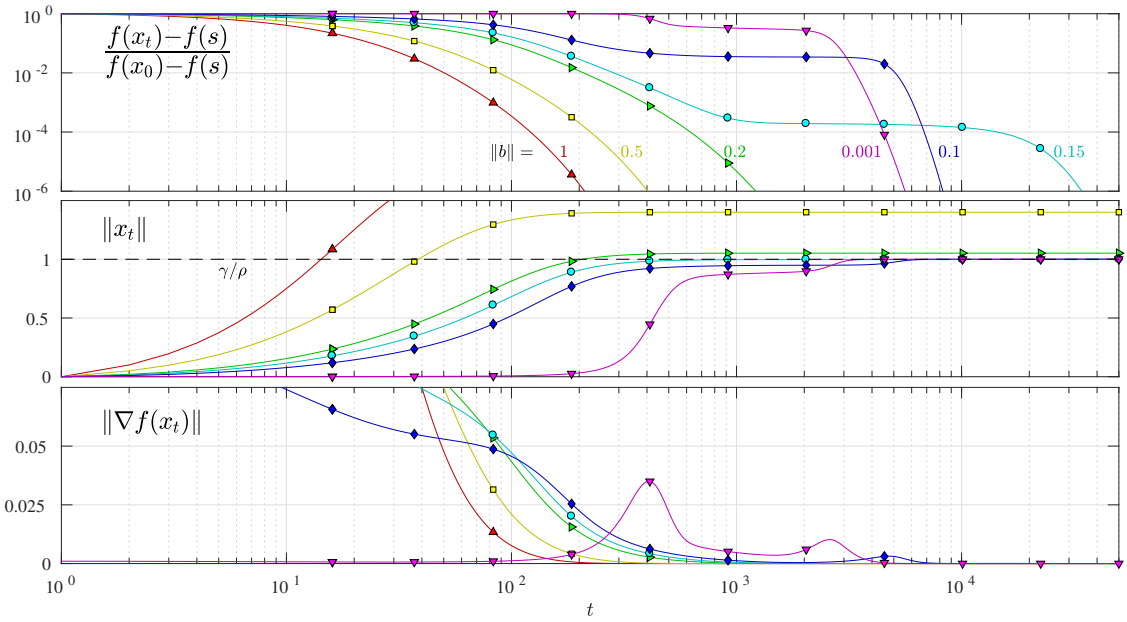

Figure 2. Trajectories of gradient descent with $\lambda^{(1)}(A) = -\gamma = -0.2$ and $\lambda^{(2)}(A), ..., \lambda^{(d)}(A)$ equally spaced between -0.18 and $\beta = 1$, and different vectors b such that $b^{(i)} = b^{(j)}$ for every $i, j \neq 1$, and $b^{(1)} = b^{(2)}/100$. The rest of the parameters are $d = 10^3$, $\eta = 0.1$, $\rho = 0.2$ and $x_0 = 0$.

(i) Corollary 3.3 highlights parallels between our guarantees and those for gradient descent on smooth convex functions [17]. In our case, $L_s/(\rho\|s\| - \gamma) \geq 1$ is a condition number, while L_s and $\|s\|$ bound the smoothness of f and iterate radius $\sup_t \|x_t\|$, respectively. We defer further comparison to Section 7.

(ii) We readily obtain relative accuracy guarantees by using the bound (4a); setting $\varepsilon = \rho\|s\|^3\varepsilon'/12$, we have $f(\tilde{x}_t) - f(s) \leq -\varepsilon' f(s) = \varepsilon'(f(0) - f(s))$, or $f(\tilde{x}_t) \leq (1-\varepsilon')f(s)$, for any $t \geq \tilde{T}_\varepsilon$, where $\tilde{T}_\varepsilon$ is defined in (11).

(iii) Evaluating $\tilde{T}_\varepsilon$ for given A, b and ρ is not straightforward, as $\|s\|$ is generally unknown. Using $\|s\| \leq R$ gives an easily computable upper bound on $\tilde{T}_\varepsilon$, and in Section 6, we demonstrate how to apply our results when $\|s\|$ is unknown.

3.2 Illustration of results

We present two experiments that investigate the behavior of gradient descent on problem (1). For the first experiment, we examine the behavior of gradient descent on single problem instances, looking at convergence behavior as we vary the vector b (to effect conditioning of the problem) by changing its norm $\|b\|$. Figure 2 summarizes our results and describes precise settings of the parameters. Recalling that $\|x_t\| \uparrow \|s\|$ (by Proposition 2.4 and Corollary 2.5), the conditioning of every instance (as given by $\rho\|s\| - \gamma$) can be read from the second plot of Figure 2, by comparing γ/ρ to $\|x_t\|$ for large t.

The plots show two behaviors of gradient descent. The problem is well-conditioned when $\|b\| \geq 0.2$, and in these cases gradient descent behaves as though the problem was strongly convex, with x_t converging linearly to s. However, for $\|b\| \leq .15$ the problem becomes ill-conditioned and gradient descent stalls around saddle points. Indeed, the third plot of Figure 2 shows that for the ill-conditioned problems, we have $\|\nabla f(x_t)\|$ increasing over some iterations, which does not occur

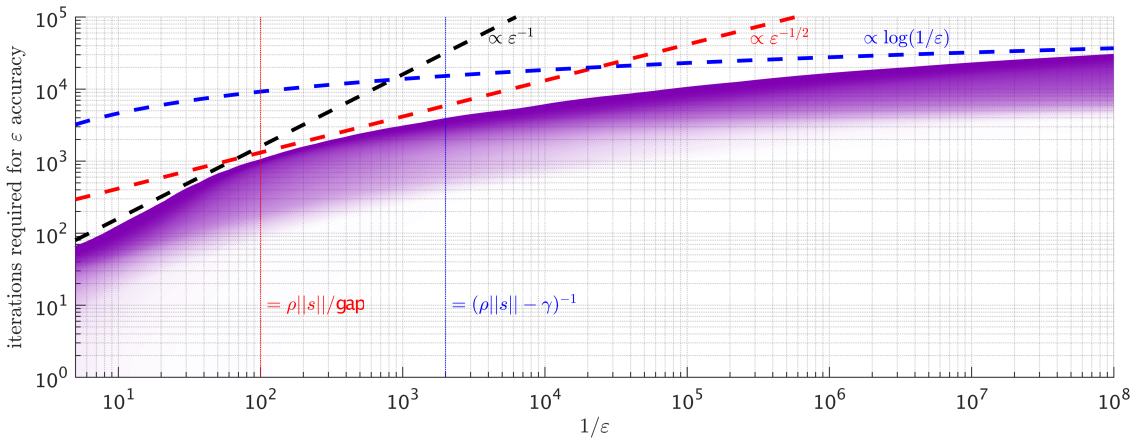

Figure 3. The purple curve is shaded according to the cdf of the number of iterations required to reach relative accuracy ε, computed over 2,500 random problem instances each with $d = 10^4$, $\beta = \rho = 1$, $\gamma = 0.5$, $\mathsf{gap} = 5 \cdot 10^{-3}$ and $\rho \|s\| - \gamma = 5 \cdot 10^{-4}$. We use $x_0 = -R_c b/\|b\|$ and $\eta = 0.25$. The black, red and blue dashed curves indicate the three convergence regimes Theorem 3.1 identifies.

in convex quadratic problems. The length of the stall does not depend only on $\rho \|s\| - \gamma$; for $\|b\| = 10^{-3}$ the stall is shorter than for $\|b\| \in \{.1, .15\}$. Instead, it appears to depend on the norm of the saddle point which causes it, which we observe from the value of $\|x_t\|$ at the time of the stall; we see that the closer the norm is to γ/ρ, the longer the stall takes. This is explained by observing that $\nabla^2 f(x) \succeq (\rho \|x\| - \gamma) I$, which means that every saddle point with norm close to γ/ρ must have only small negative curvature, and therefore harder to escape (see also Lemma 4.3 in the sequel). Fortunately, as we see in Fig. 2, saddle points with large norm have near-optimal objective value—this is the intuition behind our proof of the sub-linear convergence rates.

In our second experiment, we test our rate guarantees by considering the performance of gradient descent over an ensemble of random instances. We generate random instances with a fixed value of γ, β, ρ, $\|s\|$ and gap as follows. We set $A = \mathrm{diag}([-\gamma; -\gamma + \mathsf{gap}; u])$ with u uniformly random in $[-\gamma + \mathsf{gap}, \beta]^{d-2}$. We draw $\check{s} = (A + \rho \|s\|)^{-\zeta} \nu$, where $\nu \sim \mathcal{N}(0; I)$ and $\log_2 \zeta$ is uniform on $[-1, 1]$. We then set $s = (\|s\| / \|\check{s}\|)\check{s}$ and $b = -(A + \rho \|s\|)s$, so that s is the global minimizer of problem instance (A, b, ρ). The choice of ζ ensures we observe a large variety in the values of $\|x_t\|$ at which gradient descent stalls, allowing us to find difficult instances for each value of ε. In Figure 3 we depict the cumulative distribution of the number of iterations required to find an ε-relatively-accurate solution versus $1/\varepsilon$. The slopes in the plot agree with our upper bounds, suggesting the sharpness of our theoretical results.

4 Proofs of main results

In this section, we provide proofs of our main results, Theorems 3.1 and 3.2. A number of the steps involve technical lemmas whose proofs we defer to Appendix B. In all lemma statements, we tacitly let Assumptions A and B hold, as in the main theorem statements. Without loss of generality, we assume $\varepsilon \leq \frac{1}{2}\beta \|s\|^2 + \rho \|s\|^3$, as f is $\beta + 2\rho \|s\|$ smooth on the set $\{x : \|x\| \leq \|s\|\}$ and therefore $f(x_0) \leq f(s) + \varepsilon$ for any $\varepsilon \geq \frac{1}{2}\beta \|s\|^2 + \rho \|s\|^3$.

4.1 Proof of Theorem 3.1

We divide the proof of Theorem 3.1 into two main steps, one for each term inside the minimum in (10).

4.1.1 Linear convergence and exponential growth

We first prove that $f(x_t) \le f(s) + \varepsilon$ for $t \ge \frac{1}{\eta(\rho\|s\|-\gamma)}(\tau_{\text{grow}}(b^{(1)}) + \tau_{\text{converge}}(\varepsilon))$. We begin with two lemmas that provide regimes in which x_t converges to the solution s linearly.

Lemma 4.1. *For each $t > 0$, we have*

$$\|x_t - s\|^2 \le \left(1 - \eta\left[\rho\|x_t\| - \left(\gamma - \frac{\rho\|s\| - \gamma}{2}\right)\right]\right)\|x_{t-1} - s\|^2$$

See Appendix B.1 for a proof of this lemma.

For non-convex problem instances (those with $\gamma > 0$), the above recursion is a contraction (implying linear convergence of x_t to s) only when $\rho\|x_t\|$ is larger than $\frac{1}{2}(\rho\|s\| - \gamma)$. Using the fact that $\|x_t\|$ is non-decreasing (Corollary 2.5), Lemma 4.1 immediately implies the following result.

Lemma 4.2. *If $\rho\|x_t\| \ge \gamma - \frac{1}{2}(\rho\|s\| - \gamma) + \delta$ for some $t \ge 0$, then for all $\tau \ge 0$,*

$$\|x_{t+\tau} - s\|^2 \le (1 - \eta\delta)^\tau \|x_t - s\|^2 \le 2\|s\|^2 e^{-\eta\delta\tau}.$$

Proof. Lemma 4.1 implies that $\|x_{t+\tau} - s\|^2 \le (1 - \eta\delta)\|x_{t+\tau-1} - s\|^2$ for all $\tau > 1$. Using that $\|x_t - s\|^2 \le \|x_t\|^2 + \|s\|^2 \le 2\|s\|^2$ by Lemma 2.3, Corollary 2.5, and $1 + \alpha \le e^\alpha$ for all α gives the result. □

It remains to show that $\rho\|x_t\|$ will quickly exceed any level below γ. Fortunately, as long as $\rho\|x_t\|$ is below $\gamma - \delta$, $|x_t^{(1)}|$ grows faster than $(1 + \eta\delta)^t$:

Lemma 4.3. *Let $\delta > 0$. Then $\rho\|x_t\| \ge \gamma - \delta$ for all $t \ge \frac{2}{\eta\delta}\log(1 + \frac{\gamma_+^2}{4\rho|b^{(1)}|})$.*

See Appendix B.2 for a proof of this lemma.

We now combine the lemmas to give the linear convergence regime of Theorem 3.1. Applying Lemma 4.3 with $\delta = \frac{1}{3}(\rho\|s\| - \gamma)$ yields $\rho\|x_t\| \ge \gamma - \frac{1}{3}(\rho\|s\| - \gamma)$ for

$$t \ge T_1 \triangleq \frac{6}{\eta(\rho\|s\| - \gamma)}\log\left(1 + \frac{\gamma_+^2}{4\rho|b^{(1)}|}\right) = \frac{1}{\eta(\rho\|s\| - \gamma)}\tau_{\text{grow}}(b^{(1)}).$$

By Lemma 4.2, for any t we have

$$\|x_{T_1+t} - s\|^2 \le 2\|s\|^2 \exp\Big(-\frac{1}{6}\eta(\rho\|s\| - \gamma)\,t\Big). \tag{12}$$

As a consequence, for all $t \ge 0$ we may use the $(\beta + 2\rho\|s\|)$-smoothness of f and the fact that $\|x_t\| \le \|s\|$ (by Corollary 2.5) to obtain

$$f(x_t) - f(s) \le \frac{\beta + 2\rho\|s\|}{2}\|x_t - s\|^2 \le (\beta + 2\rho\|s\|)\|s\|^2 e^{-\frac{1}{6}\eta(\rho\|s\|-\gamma)(t - T_1)}$$

where we have used that $\nabla f(s) = 0$ and the bound (12). Therefore, if we set

$$T_2 \triangleq \frac{6}{\eta(\rho\|s\| - \gamma)}\log\frac{(\beta + 2\rho\|s\|)\|s\|^2}{\varepsilon} = \frac{1}{\eta(\rho\|s\| - \gamma)}\tau_{\text{converge}}(\varepsilon),$$

then $t \ge T_1 + T_2 = \frac{1}{\eta(\rho\|s\|-\gamma)}(\tau_{\text{grow}}(b^{(1)}) + \tau_{\text{converge}}(\varepsilon))$ implies $f(x_t) - f(s) \le \varepsilon$.

4.1.2 Sublinear convergence and convergence in subspaces

We now turn to the sublinear convergence regime in Theorem 3.1, which applies when the quantity $\rho\|s\| - \gamma$ is so small as to cause the guarantees of Section 4.1.1 to be irrelevant. Consider the condition

$$\rho\|s\| - \gamma \le \frac{\varepsilon}{10\|s\|^2}\ . \tag{13}$$

Note that if (13) fails to hold, then (10) is dominated by the $(\rho\|s\| - \gamma)^{-1}$ term. Therefore, to complete the proof of Theorem 3.1 it suffices to show that if (13) holds, then $f(x_t) \le f(s) + \varepsilon$ whenever

$$t \ge T_\varepsilon^{\text{sub}} \triangleq \frac{\tau_{\text{grow}}(b^{(1)}) + \tau_{\text{converge}}(\varepsilon)}{\eta} \min\left\{\frac{10\|s\|^2}{\varepsilon}, \sqrt{\frac{10\|s\|^2}{(\mathsf{gap} \wedge \rho\|s\|)\varepsilon}}\right\}. \tag{14}$$

Roughly, our proof of the result (14) proceeds as follows: when $\rho\|s\| - \gamma$ is small, the function f is very smooth along eigenvectors with eigenvalues close to $-\gamma = \lambda^{(1)}(A)$. It is therefore sufficient to show convergence in the complementary subspace, which occurs at a linear rate. Appropriately choosing the gap between the eigenvalues in the complementary subspace and $\lambda^{(1)}(A)$ to trade between convergence rate and function smoothness yields the rates (14).

The following analogs of Lemmas 4.1 and 4.2 establish subspace convergence.

Lemma 4.4. *Let Π be any projection matrix satisfying $\Pi A = A\Pi$ for which $\Pi A_s \succeq \delta\Pi$ for some $\delta > 0$. For all $t > 0$,*

$$\begin{aligned}\|\Pi A_s^{1/2}(x_t - s)\|^2 &\le (1 - \eta\delta)\|\Pi A_s^{1/2}(x_{t-1} - s)\|^2 \\ &\quad + \sqrt{8}\eta\rho\left(\|s\| - \|x_{t-1}\|\right)\left[\rho\left(\|s\| - \|x_{t-1}\|\right)\|x_{t-1}\|^2 + \|(I - \Pi)A_s\|_{\text{op}}\|s\|^2\right].\end{aligned}$$

See Appendix B.3 for a proof. Letting $\Pi_\delta = \sum_{i:\lambda^{(i)} \ge \delta + \lambda^{(1)}} v_i v_i^T$ be the projection matrix onto the span of eigenvectors of A with eigenvalues at least $\lambda^{(1)}(A) + \delta$, we obtain the following consequence of Lemma 4.4, whose proof we provide in Appendix B.4.

Lemma 4.5. *Let $t \ge 0$, $\delta \ge 0$, and define $\bar\delta = \delta \vee \mathsf{gap}$. If $\rho\|s\| \le \gamma + \sqrt{\delta\bar\delta}$ and $\rho\|x_t\| \ge \gamma - \frac{1}{3}\sqrt{\delta\bar\delta}$, then for any $\tau \ge 0$,*

$$\begin{aligned}\left\|\Pi_\delta A_s^{1/2}(x_{t+\tau} - s)\right\|^2 &\le \left(1 - \eta\bar\delta\right)^\tau \left\|\Pi_\delta A_s^{1/2}(x_t - s)\right\|^2 + 13\|s\|^2\delta \\ &\le 2\left(\beta + \rho\|s\|\right)\|s\|^2 e^{-\eta\bar\delta\tau} + 13\|s\|^2\delta.\end{aligned}$$

We use these lemmas to prove the desired bound (14) by appropriate separation of the eigenspaces over which we guarantee convergence. To that end, we define

$$\delta \triangleq \frac{\varepsilon}{10\|s\|^2}\ \ , \ \ \bar\delta \triangleq \delta \vee \mathsf{gap} \ \text{ and } \ \bar\delta' \triangleq \delta \vee (\mathsf{gap} \wedge \rho\|s\|) \le \bar\delta, \tag{15}$$

and note that the definition of gap immediately implies $\Pi_\delta = \Pi_{\bar\delta}$. The growth guranteed by Lemma 4.3 shows that $\rho\|x_t\| \ge \gamma - \frac{1}{3}\sqrt{\delta\bar\delta'}$ for every

$$t \ge T_1^{\text{sub}} \triangleq \frac{6}{\eta\sqrt{\delta\bar\delta'}}\log\left(1 + \frac{\gamma_+^2}{4\rho|b^{(1)}|}\right) = \frac{1}{\eta\sqrt{\delta\bar\delta'}}\tau_{\text{grow}}(b^{(1)}).$$

Additionally, for $t \geq T_1^{\text{sub}}$ we have $\rho \|x_t\| \geq \gamma - \frac{1}{3}\sqrt{\delta\bar{\delta}'} \geq \gamma - \frac{1}{3}\sqrt{\delta\bar{\delta}}$ because $\bar{\delta} \geq \bar{\delta}'$. Thus, if we define[1]

$$T_2^{\text{sub}} \triangleq \frac{1}{\eta\bar{\delta}} \log \frac{2(\beta + \rho\|s\|)}{\delta} \leq \frac{1}{\eta\sqrt{\delta\bar{\delta}'}} \log \left(\left[\frac{(\beta + 2\rho\|s\|)\|s\|^2}{\varepsilon} \right]^6 \right) = \frac{\tau_{\text{converge}}(\varepsilon)}{\eta\sqrt{\delta\bar{\delta}'}}$$

we may use Lemma 4.5 to find that for all $t \geq T^{\text{sub}} = T_1^{\text{sub}} + T_2^{\text{sub}}$, we have

$$\left\| \Pi_\delta A_s^{1/2}(x_t - s) \right\|^2 \leq 14\|s\|^2 \delta. \tag{16}$$

We now translate the guarantee (16) on the distance from x_t to s in the subspace of "large" eigenvectors of A to a guarantee on the solution quality $f(x_t)$. Using the expression (5) for $f(x)$, the orthogonality of $I - \Pi_\delta$ and Π_δ and $\|x_t\| \leq \|s\|$, we have

$$f(x_t) \leq f(s) + \frac{1}{2}\left\| (I - \Pi_\delta) A_s^{1/2}(x_t - s) \right\|^2 + \frac{1}{2}\left\| \Pi_\delta A_s^{1/2}(x_t - s) \right\|^2 + \frac{\rho\|s\|}{2}(\|s\| - \|x_t\|)^2.$$

Now we note that

$$\|(I - \Pi_\delta)A_s\|_{\text{op}} = \max_{i:\lambda^{(i)} < \lambda^{(1)} + \delta} |\lambda^{(i)} + \rho\|s\|| \leq -\gamma + \delta + \rho\|s\| \leq 2\delta,$$

where we have used our assumption (13) that $\rho\|s\| - \gamma \leq \frac{\varepsilon}{10\|s\|^2} = \delta$. Using this gives

$$f(x_t) \leq f(s) + \delta\|x_t - s\|^2 + 7\|s\|^2\delta + \frac{\rho\|s\|}{2}(\|s\| - \|x_t\|)^2, \tag{17}$$

where we use inequality (16). Because $\rho\|x_t\| \geq \gamma - \frac{1}{3}\sqrt{\delta\bar{\delta}'}$ for $t \geq T_1^{\text{sub}}$, we obtain

$$0 \leq \rho(\|s\| - \|x_t\|) \leq \rho\|s\| - \gamma - (\rho\|x_t\| - \gamma) \leq \frac{4}{3}\sqrt{\delta\bar{\delta}'}.$$

The above inequality provides an upper bound on $(\|s\| - \|x_t\|)^2$. Alternatively, we may bound $(\|s\| - \|x_t\|)^2 \leq \|s\|^2$ using $\|x_t\| \leq \|s\|$ (Corollary 2.5). Therefore

$$\frac{\rho\|s\|}{2}(\|s\| - \|x_t\|)^2 \leq \|s\|^2 \min\left\{ \frac{\rho\|s\|}{2}, \frac{16\bar{\delta}'\delta}{18\rho\|s\|} \right\} \leq \|s\|^2\delta,$$

where the final inequality follows from $\bar{\delta}' \leq \delta \vee \rho\|s\|$. Substituting the above display into (17) with $\|s - x_t\|^2 \leq 2\|s\|^2$ (by Lemma 2.3), we find

$$f(x_t) \leq f(x_t) + 9\|s\|^2\delta \leq f(x_t) + 10\|s\|^2\delta = f(x_t) + \varepsilon,$$

where in the equality we substitute $\delta = \frac{\varepsilon}{10\|s\|^2}$. Summarizing, we see that on the event $\rho\|s\| - \gamma \leq \delta = \frac{\varepsilon}{10\|s\|^2}$, the point x_t is ε-suboptimal for problem (1) whenever $t \geq \frac{\tau_{\text{grow}}(b^{(1)}) + \tau_{\text{converge}}(\varepsilon)}{\eta\sqrt{\delta\bar{\delta}'}} \geq T_1^{\text{sub}} + T_2^{\text{sub}}$, where $\sqrt{\delta\bar{\delta}'} = \max\{\frac{\varepsilon}{10\|s\|^2}, \sqrt{\frac{\varepsilon}{10\|s\|^2} \cdot (\mathsf{gap} \wedge \rho\|s\|)}\}$.

[1] the inequality follows from $\delta, \bar{\delta}' \leq \bar{\delta}$ and $(\beta + 2\rho\|s\|)\|s\|^2/\varepsilon \geq 2$ (see beginning of Section 4).

4.2 Proof of Theorem 3.2

Theorem 3.2 follows from three basic observations about the effect of adding a small uniform perturbation to b, which we summarize in the following lemma (see Section B.5 for a proof).

Lemma 4.6. *Set $\tilde{b} = b + \sigma q$, where q is uniform on the unit sphere in $\mathbb{R}^d$ and $\sigma > 0$. Let $\tilde{f}(x) = \frac{1}{2}x^T Ax + \tilde{b}^T x + \frac{1}{3}\rho \|x\|^3$ and let $\tilde{s}$ be a global minimizer of $\tilde{f}$. Then*

(i) For $d > 2$, $\mathbb{P}(|\tilde{b}^{(1)}| \le \sqrt{\pi}\sigma\delta/\sqrt{2d}) \le \delta$

(ii) $|f(x) - \tilde{f}(x)| \le \sigma \|x\|$ for all $x \in \mathbb{R}^d$

(iii) $\big|\, \|s\|^2 - \|\tilde{s}\|^2 \,\big| \le 2\sigma/\rho$.

With Lemma 4.6 in hand, our proof proceeds in three parts: in the first two, we provide bounds on the iteration complexity of each of the modes of convergence that Theorem 3.1 exhibits in the perturbed problem with vector $\tilde{b}$. The final part shows that the quality of the (approximate) solutions $\tilde{x}_t$ and $\tilde{s}$ is not much worse than s.

Let $\tilde{f}, \tilde{b}$ and $\tilde{s}$ be as defined in Lemma 4.6. By Theorem 3.1, we know that $\tilde{f}(\tilde{x}_t) \le \tilde{f}(\tilde{s}) + \varepsilon$ for all

$$t \ge \frac{6}{\eta}\left(\log\left(1 + \frac{\gamma_+^2/4}{\rho|\tilde{b}^{(1)}|}\right) + \log\frac{(\beta + 2\rho\|\tilde{s}\|)\|\tilde{s}\|^2}{\varepsilon}\right)\min\left\{\frac{1}{\rho\|\tilde{s}\| - \gamma}, \frac{10\|\tilde{s}\|^2}{\varepsilon}\right\}, \tag{18a}$$

and that if $\rho\|\tilde{s}\| - \gamma \le \frac{\varepsilon}{10\|\tilde{s}\|}$, then $\tilde{f}(\tilde{x}_t) \le \tilde{f}(\tilde{s}) + \varepsilon$ for all

$$t \ge \frac{6}{\eta}\left(\log\left(1 + \frac{\gamma_+^2}{4\rho|\tilde{b}^{(1)}|}\right) + \log\frac{(\beta + 2\rho\|\tilde{s}\|)\|\tilde{s}\|^2}{\varepsilon}\right)\sqrt{\frac{10\|\tilde{s}\|^2}{\varepsilon} \cdot \frac{1}{\mathsf{gap} \wedge \rho\|\tilde{s}\|}}. \tag{18b}$$

We now turn to bounding expressions (18a) and (18b) appropriately.

Part 1: upper bounding terms outside the log. Recalling that $\sigma = \frac{\rho\overline{\sigma}\varepsilon}{12(\beta+2\rho\|s\|)}$ and $\varepsilon \le (\frac{1}{2}\beta + \rho\|s\|)\|s\|^2$, we have $\sigma \le \frac{\rho}{24}\overline{\sigma}\|s\|^2$. Thus, part (iii) of Lemma 4.6 gives

$$|\, \|s\|^2 - \|\tilde{s}\|^2 \,| \le 2\sigma/\rho \le \overline{\sigma}\|s\|^2/12, \quad \text{so} \quad \|\tilde{s}\|^2 \in (1 \pm \overline{\sigma}/12)\|s\|^2.$$

Consequently, using $\overline{\sigma} \le 1$ we have

$$\big|\, \|s\| - \|\tilde{s}\| \,\big| \le \frac{2\sigma}{\rho(\|s\| + \|\tilde{s}\|)} \le \frac{2\overline{\sigma}\varepsilon}{12(1 + \sqrt{11/12})\|s\|(\beta + 2\rho\|s\|)} \le \frac{\overline{\sigma}\varepsilon}{20\rho\|s\|^2}. \tag{19}$$

Now, suppose that $\frac{\varepsilon}{10\|s\|^2} \le \rho\|s\| - \gamma$. Substituting this into the bound (19) yields $|\, \|s\| - \|\tilde{s}\| \,| \le \frac{\overline{\sigma}}{2\rho}(\rho\|s\| - \gamma)$, and rearranging, we obtain

$$\rho\|\tilde{s}\| - \gamma \ge (1 - 0.5\overline{\sigma})(\rho\|s\| - \gamma) \ge \frac{\rho\|s\| - \gamma}{1 + \overline{\sigma}}$$

because $\overline{\sigma} \le 1$. We combine the preceding bounds to obtain

$$\min\left\{\frac{1}{\rho\|\tilde{s}\| - \gamma}, \frac{10\|\tilde{s}\|^2}{\varepsilon}\right\} \le (1 + \overline{\sigma})\min\left\{\frac{1}{\rho\|s\| - \gamma}, \frac{10\|s\|^2}{\varepsilon}\right\} \tag{20a}$$

and

$$\sqrt{\frac{10\left\|\tilde{s}\right\|^2}{\varepsilon}\cdot\frac{1}{\mathsf{gap}\wedge\rho\left\|\tilde{s}\right\|}}\leq(1+\overline{\sigma})\sqrt{\frac{10\left\|s\right\|^2}{\varepsilon}\cdot\frac{1}{\mathsf{gap}\wedge\rho\left\|s\right\|}}\,, \tag{20b}$$

where we have used $\left\|\tilde{s}\right\|\leq(1+\overline{\sigma})\left\|s\right\|^2$ and $\left\|\tilde{s}\right\|\geq\sqrt{1-\overline{\sigma}/12}\left\|s\right\|^2\geq\left\|s\right\|^2/(1+\overline{\sigma})$.

The bound (19) also implies $\rho\left\|\tilde{s}\right\|-\gamma\leq\rho\left\|s\right\|-\gamma+\frac{\overline{\sigma}\varepsilon}{20\|s\|^2}$. When $\rho\left\|s\right\|-\gamma\leq(1-2\overline{\sigma}/3)\frac{\varepsilon}{10\|s\|^2}$, we thus have

$$\rho\left\|\tilde{s}\right\|-\gamma\leq\frac{\varepsilon}{10\left\|s\right\|^2}\left(1-\frac{2}{3}\overline{\sigma}+\frac{1}{2}\overline{\sigma}\right)\leq\frac{\varepsilon(1+\overline{\sigma}/12)(1-\overline{\sigma}/6)}{10\left\|\tilde{s}\right\|^2}\leq\frac{\varepsilon}{10\left\|\tilde{s}\right\|^2},$$

where we have used $\left\|\tilde{s}\right\|^2\leq(1+\overline{\sigma}/12)\left\|s\right\|^2$. The condition $\rho\left\|s\right\|-\gamma\leq(1-2\overline{\sigma}/3)\frac{\varepsilon}{10\|s\|^2}$ thus implies that the bound (18b) is in force.

Part 2: upper bounding terms inside the log. Fix a confidence level $\delta\in(0,1)$. By Lemma 4.6(i), $1/|\tilde{b}^{(1)}|\leq\sqrt{2d}/(\sqrt{\pi}\sigma\delta)\leq\sqrt{d}/(\sigma\delta)$ with probability at least $1-\delta$, so

$$\begin{aligned}6\log\left(1+\frac{\gamma_+^2}{4\rho|\tilde{b}^{(1)}|}\right)\leq 6\log\left(1+\frac{\gamma_+^2\sqrt{d}}{4\rho\sigma\delta}\right)&\overset{(\star)}{\leq}6\log\left(1+\mathbb{I}_{\{\gamma>0\}}\frac{3\sqrt{d}}{\overline{\sigma}\delta}\right)+\\&6\log\frac{(\beta+2\rho\left\|s\right\|)\left\|s\right\|^2}{\varepsilon}=\tilde{\tau}_{\mathrm{grow}}(d,\delta,\overline{\sigma})+\frac{6}{14}\tilde{\tau}_{\mathrm{converge}}(\varepsilon),\end{aligned}$$

where inequality $(\star)$ uses that $\rho\left\|s\right\|\geq\gamma_+$ and $\varepsilon\leq(\beta+\frac{1}{2}\rho\left\|s\right\|)\left\|s\right\|^2$. Using $\left\|\tilde{s}\right\|\leq\sqrt{1+\overline{\sigma}/12}\left\|s\right\|$ yields the upper bound

$$6\log\frac{(\beta+2\rho\left\|\tilde{s}\right\|)\left\|\tilde{s}\right\|^2}{\varepsilon}\leq 6\log\frac{(\beta+2\rho\left\|s\right\|)\left\|s\right\|^2}{\varepsilon}+9\log(1+\overline{\sigma}/12)\leq\frac{8}{14}\tilde{\tau}_{\mathrm{converge}}(\varepsilon),$$

where the second inequality follows from $9\log(1+\overline{\sigma}/12)<2\log 2\leq 2\log\frac{(\beta+2\rho\|s\|)\|s\|^2}{\varepsilon}$.

Substituting the above bounds and the upper bounds (20a) and (20b) into expressions (18a) and (18b), we see that the iteration bounds claimed in Theorem 3.2 hold. To complete the proof we need only bound the quality of the solution $\tilde{x}_t$.

Part 3: bounding solution quality. We recall that $\sigma=\frac{\rho\overline{\sigma}\varepsilon}{12(\beta+2\rho\|s\|)}\leq\frac{\overline{\sigma}\varepsilon}{24\|s\|}$ and $\left\|\tilde{s}\right\|\leq\sqrt{1+\overline{\sigma}/12}\left\|s\right\|\leq\sqrt{2}\left\|s\right\|$, so $\sigma\leq\frac{\overline{\sigma}\varepsilon}{\|s\|+\|\tilde{s}\|}$. Thus, whenever $\tilde{f}(\tilde{x}_t)\leq\tilde{f}(\tilde{s})+\varepsilon$,

$$\begin{aligned}f(\tilde{x}_t)&\overset{(a)}{\leq}\tilde{f}(\tilde{x}_t)+\sigma\left\|\tilde{x}_t\right\|\leq\tilde{f}(\tilde{s})+\varepsilon+\sigma\left\|\tilde{x}_t\right\|\overset{(b)}{\leq}\tilde{f}(\tilde{s})+\varepsilon+\sigma\left\|\tilde{s}\right\|\\&\overset{(c)}{\leq}\tilde{f}(s)+\varepsilon+\sigma\left\|\tilde{s}\right\|\overset{(d)}{\leq}f(s)+\sigma(\left\|\tilde{s}\right\|+\left\|s\right\|)+\varepsilon\leq f(s)+(1+\overline{\sigma})\varepsilon,\end{aligned}$$

where transitions (a) and (d) follow from part (ii) of Lemma 4.6, transition (b) follows from $\left\|\tilde{x}_t\right\|\leq\left\|\tilde{s}\right\|$ (Corollary 2.5), and transition (c) follows from $\tilde{f}(\tilde{s})=\min_{z\in\mathbb{R}^d}\tilde{f}(z)$.

5 Convergence of a line search method

The maximum step size allowed by Assumption A may be too conservative (as is frequent with gradient descent). With that in mind, in this section we briefly analyze line search schemes of the form

$$x_{t+1}=x_t-\eta_t\nabla f(x_t)\quad\text{where}\quad\eta_t=\operatorname*{argmin}_{\eta\in C_t}f(x_t-\eta\nabla f(x_t)) \tag{21}$$

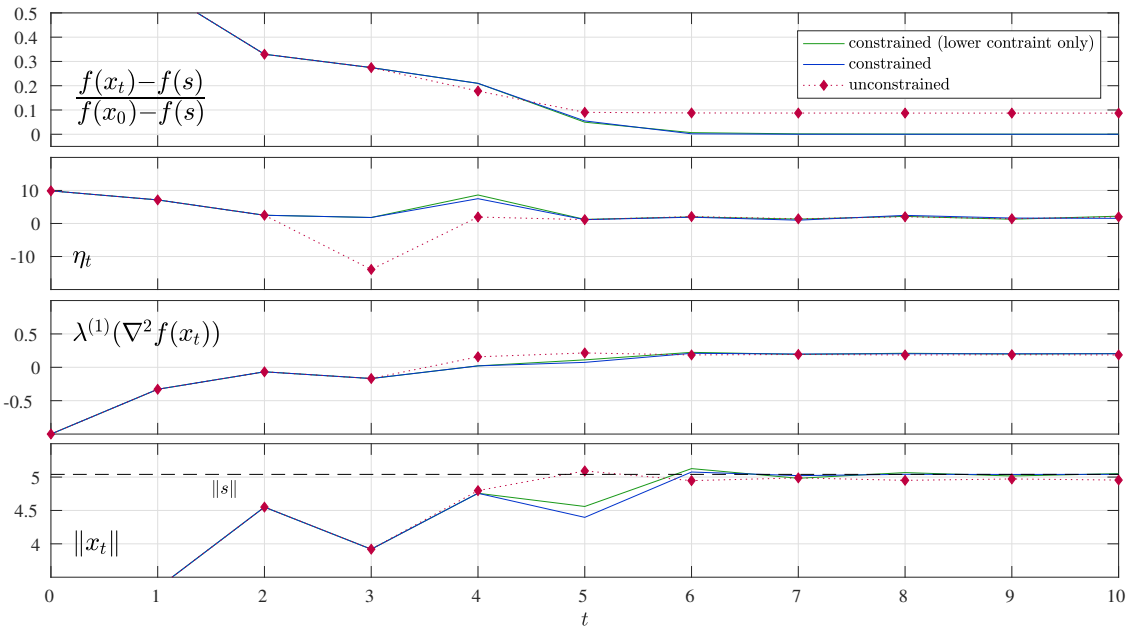

Figure 4. Steepest descent variants applied on $A = \mathrm{diag}([-1; -0.8; -0.5])$, $b = [0.04; 0.15; 0.3]$ and $\rho = 0.2$. The red, green, and blue curves correspond to $C_t = \mathbb{R}$, $C_t = [0, \infty)$ and C_t given by (22), respectively.

and C_t is a (possibly time-varying) interval of allowed step sizes. For the problem (1), η_t is computable for any interval C_t, as the critical points of the function $h(\eta) = f(x_t - \eta \nabla f(x_t))$ are roots of a quatric polynomial with coefficients determined by $\|x\|, \|g\|, g^T A g$, and $x^T g$, so η_t must be a root or an edge of the interval C_t.

The unconstrained choice $C_t = \mathbb{R}$ yields the *steepest descent* method [19]. As we demonstrate presently, for steepest descent it is possible that $\eta_t < 0$ and that convergence to a sub-optimal local minimum of f occurs. Consequently, we propose choosing the updates (21) using the interval

$$C_t = \left[0, \left[\frac{\nabla f(x_t)^T A \nabla f(x_t)}{\|\nabla f(x_t)\|^2} + \rho \|x_t\|\right]_+^{-1}\right]. \tag{22}$$

The scheme (22) is guaranteed to converge to the global minimum of f, as we formalize in the following proposition, whose proof we provide in Appendix C.

Proposition 5.1. *Let x_t be the iterates of gradient descent with step sizes selected by the constrained minimization* (22)*. Let Assumption B hold and assume $b^{(1)} \neq 0$. Then s is the unique global minimizer of f and $\lim_{t\to\infty} x_t = s$.*

In Fig. 4, we display the quantities $f(x_t), \eta_t, \lambda^{(1)}(\nabla^2 f(x_t))$, and $\|x_t\|$ for the above line-search variants on a $d = 3$-dimensional problem instance. The step sizes differ at iteration $t = 3$, where the unconstrained gradient step makes almost 50% more progress than steps restricted to be positive. However, it then converges to a sub-optimal local minimum (note $\lambda^{(1)}(\nabla^2 f(x_t)) > 0$) approximately 9% worse than the global minimum achieved by the guarded sequence (22). The step sizes these methods choose are significantly larger than the η Assumption A allows, which is approximately 0.12. Fig. 4 reveals a difference between fixed step size gradient descent and the line-search schemes—the norm $\|x_t\|$ of the line-search-based iterates is non-monotonic and overshoots $\|s\|$. Our convergence

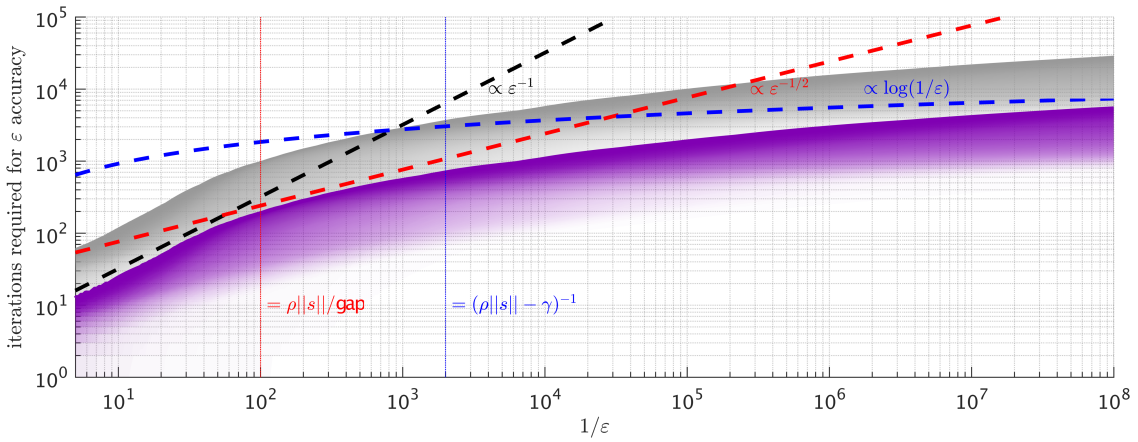

Figure 5. Reproduction of of the ensemble experiment reported in Section 3.2 (Figure 3), with the scheme (22) used instead of fixed step size gradient descent. The cdf for fixed step size $\eta = 0.25$ is shown in gray for comparison.

rate proofs hinge on Corollary 2.5, that $\|x_t\|$ is increasing, so extension of our rates to line-search schemes is not straightforward.

We believe that the rate guarantees of Theorem 3.1 apply also to the step size choice (22). To lend credence to this hypothesis, we repeat the ensemble experiment detailed in Section 3.2 (Figure 3), where we use the step size (22) instead of the fixed step size. Figure 5 shows that the rates we prove in Section 3 seem to accurately describe the behavior of guarded steepest descent as well, with constant factors.

We remark that we introduce the upper constraint (22) only because we require it in the proof of Proposition 5.1. Empirically, a scheme with the simpler constraint $C_t = [0, \infty)$ appears to converge to the global minimum as well, though we remain unable to prove this. While such step size can differ from the choice (22) (see time $t = 4$ in Fig. 4), the variants seem equally practicable. Indeed, we performed the ensemble experiment (Figs. 3 and 5) with $C_t = [0, \infty)$ and the results are indistinguishable.

6 Application: A Hessian-free majorization method

As one of our motivations is in the application of cubic-regularized problems of the form (1) to general (smooth) nonlinear optimization problems, in this section we develop an optimization method that uses gradient descent to approximately solve the Nesterov-Polyak cubic regularized Newton step (2). We consider functions g satisfying the following

Assumption C. *The function g satisfies* $\inf g \geq g^\star > -\infty$, *is β-smooth and has 2ρ-Lipschitz Hessian,* i.e. $\|\nabla^2 g(y) - \nabla^2 g(y')\| \leq 2\rho\|y - y'\|$ *for every $y, y' \in \mathbb{R}^d$.*

The first two parts—boundedness and smoothness—of Assumption C are standard. The third implies that g is upper bounded by its cubic-regularized quadratic approximation [18, Lemma 1]: for $y, \Delta \in \mathbb{R}^d$ one has

$$g(y + \Delta) \leq g(y) + \nabla g(y)^T \Delta + \frac{1}{2}\Delta^T \nabla^2 g(y)\Delta + \frac{\rho}{3}\|\Delta\|^3 . \tag{23}$$

Following [18, 1, 3], our goal is to find an ϵ-second-order stationary point y_ϵ:

$$\|\nabla g(y_\epsilon)\| \leq \epsilon \quad \text{and} \quad \nabla^2 g(y_\epsilon) \succeq -\sqrt{\rho\epsilon}\, I. \tag{24}$$

Algorithm 1 A second-order majorization method

1: **function** SOLVE-PROBLEM(y_0, g, β, ρ, ϵ, δ)
2: Set $K_{\text{prog}} = 1/324$, $\eta = 1/(10\beta)$ and $K_{\max} = \lceil (g(y_0) - g^\star)\sqrt{\rho}/(K_{\text{prog}}\epsilon^{3/2}) \rceil$
3: Set $\delta' = \delta / K_{\max}$
4: **for** $k = 1, 2, \ldots$ **do** ▷ guaranteed to return at $k \le K_{\max}$
5: $\Delta_k \leftarrow$ SOLVE-SUBPROBLEM($\nabla^2 g(y_{k-1})$, $\nabla g(y_{k-1})$, ρ, η, $\sqrt{\epsilon/(9\rho)}$, $1/2$, δ')
6: **if** $g(y_{k-1} + \Delta_k) \le g(y_k) - K_{\text{prog}}\epsilon^{3/2}\rho^{-1/2}$ **then**
7: $y_k \leftarrow y_{k-1} + \Delta_k$
8: **else**
9: $\Delta_k \leftarrow$ SOLVE-FINAL-SUBPROBLEM($\nabla^2 g(y_{k-1})$, $\nabla g(y_{k-1})$, ρ, η, $\epsilon/2$)
10: **return** $y_{k-1} + \Delta_k$

Intuitively, ϵ-second-order stationary points provide a finer approximation to local minima than ϵ-stationary points (with only $\|\nabla g(y)\| \le \epsilon$). Throughout this section, we use ϵ to denote approximate stationarity in g, and continue to use ε to denote approximate optimality for subproblems of the form (1).

We outline a majorization-minimization strategy [6, 19] for optimization of g in Algorithm 1. At each iteration, the method (approximately) minimizes a local model of g, halting once progress decreasing g falls below a certain threshold. In Algorithm 2, we describe our Hessian-free approximate solution of the subproblem (23) using gradient descent with a small perturbation to the linear term; we write the method in terms of an input matrix $A = \nabla^2 g(y)$, noting that it requires only matrix-vector products Av implementable by a first-order oracle for g.

The method SOLVE-SUBPROBLEM takes as input a problem instance (A, b, ρ), confidence level δ, relative accuracy ε', and a putative lower bound on the magnitude of the global minimizer s, which we denote by $\|\hat{s}\|$. As an immediate consequence of Theorem 3.2, as long as $\|s\| \ge \|\hat{s}\|$ the method is guaranteed to terminate before reaching line 10, and if the gradient is sufficiently large, termination occurs before entering the loop. We formalize this in the following lemma, whose proof we provide in Appendix D.1.

Lemma 6.1. *Let $A \in \mathbb{R}^{d\times d}$ satisfy $\|A\|_{\text{op}} \le \beta$, $b \in \mathbb{R}^d$, $\rho > 0, \|\hat{s}\| > 0$, $\varepsilon' \in (0,1), \delta \in (0,1)$ and $\eta \le 1/(8\beta + 4\rho\|\hat{s}\|)$. With probability at least $1 - \delta$, if*

$$\|s\| \ge \|\hat{s}\| \text{ or } \|b\| \ge \max\{\sqrt{\beta\rho}\|\hat{s}\|^{3/2}, \rho\|\hat{s}\|^2\}$$

then $x =$ SOLVE-SUBPROBLEM$(A, b, \rho, \eta, \|\hat{s}\|, \varepsilon', \delta)$ satisfies $f(x) \le (1 - \varepsilon')\rho\|\hat{s}\|^3/6$.

Let s_k be the global minimizer (in Δ) of the model (23) at $y = y_k$, the kth iterate of Algorithm 1. Lemma 6.1 guarantees that with high probability, if SOLVE-SUBPROBLEM fails to meet the progress condition in line 6 at iteration k, then $\|s_k\| \le \sqrt{\epsilon/(9\rho)}$, and therefore $\lambda^{(1)}(\nabla^2 g(y_k)) \ge -\rho\|s_k\| \ge -\sqrt{\rho\epsilon}$. It is possible, nonetheless, that $\|\nabla g(y_k)\| > \epsilon$; to address this, we correctively apply gradient descent on the final subproblem (SOLVE-FINAL-SUBPROBLEM).

Building off of an argument of Nesterov and Polyak [18, Lemma 5], we obtain the following guarantee for Algorithm 1, whose proof we provide in Appendix D.2.

Proposition 6.2. *Let g satisfy Assumption C, $y_0 \in \mathbb{R}^d$ be arbitrary, and let $\delta \in (0, 1]$ and $\epsilon \le \min\{\beta^2/\rho, \rho^{1/3}\}$. With probability at least $1 - \delta$, Algorithm 1 finds an ϵ-second-order stationary point (24) in at most*

$$O(1) \cdot \frac{\beta\,\lceil g(y_0) - g^\star \rceil}{\epsilon^2} \log\left(\frac{d}{\delta} \cdot \frac{\beta\,\lceil g(y_0) - g^\star \rceil}{\epsilon^2}\right) \tag{25}$$

Algorithm 2 A Hessian-free subproblem solver

function SOLVE-SUBPROBLEM(A, b, ρ, η, $\|\hat{s}\|$, ε', δ)
  Set $f(x) = (1/2)x^T Ax + b^T x + (\rho/3)\|x\|^3$ and $x_0 = \text{CAUCHY-POINT}(A, b, \rho)$
  if $f(x_0) \leq -(1-\varepsilon')\rho\|\hat{s}\|^3/6$ **then return** x_0
  Set

$$T = \frac{240}{\eta\rho\|\hat{s}\|\varepsilon'}\left[6\log\left(1+\frac{\sqrt{d}}{\delta}\right) + 32\log\left(\frac{6}{\eta\rho\|\hat{s}\|\varepsilon'}\right)\right]$$

  Set $\sigma = \frac{\rho^2\|\hat{s}\|^3\varepsilon'}{144(\beta+2\rho\|\hat{s}\|)}$, draw q uniformly from the unit sphere, set $\tilde{b} = b + \sigma q$
  Set $\tilde{f}(x) = (1/2)x^T Ax + \tilde{b}^T x + (\rho/3)\|x\|^3$ and $\tilde{x}_0 = \text{CAUCHY-POINT}(A, \tilde{b}, \rho)$
  for $t = 1, 2, \ldots, T$ **do**
    $\tilde{x}_t \leftarrow \tilde{x}_{t-1} - \eta\nabla\tilde{f}(\tilde{x}_{t-1})$
    if $f(\tilde{x}_t) \leq -(1-\varepsilon')\rho\|\hat{s}\|^3/6$ **then return** $\tilde{x}_t$
  return $\tilde{x}_t$

function SOLVE-FINAL-SUBPROBLEM(A, b, ρ, η, ε_{g})
  Set $f(x) = (1/2)x^T Ax + b^T x + (\rho/3)\|x\|^3$ and $\Delta = \text{CAUCHY-POINT}(A, b, \rho)$
  while $\|\nabla f(\Delta)\| > \varepsilon_{\mathrm{g}}$ **do** $\Delta \leftarrow \Delta - \eta\nabla f(\Delta)$
  return Δ

function CAUCHY-POINT(A, b, ρ)
  return $-R_c b/\|b\|$ where $R_c = \frac{-b^T Ab}{2\rho\|b\|^2} + \sqrt{\left(\frac{b^T Ab}{2\rho\|b\|^2}\right)^2 + \frac{\|b\|}{\rho}}$

gradient and Hessian-vector product evaluations.

In Proposition 6.2, the assumption $\epsilon \leq \beta^2/\rho$ is no loss of generality, as otherwise the Hessian guarantee (24) is trivial, and we may obtain the gradient guarantee by simply running gradient descent on g for $2\beta(g(y_0) - g^\star)\epsilon^{-2}$ iterations. Similarly, if $\epsilon > \rho^{1/3}$ then $K_{\max} \leq 324\lceil g(y_0) - g^\star\rceil$ and the proof of Proposition 6.2 reveals that the overall iteration count scales as $\epsilon^{-1/2}$ instead of ϵ^{-2}.

There are other Hessian-free methods that provide the guarantee (24), and recent schemes using acceleration techniques [1, 3] provide it in roughly $\epsilon^{-7/4}\log\frac{d}{\delta}$ first-order operations, which is better than Algorithm 1. Nevertheless, this section illustrates how gradient descent on the structured problem (1) can be straightforwardly leveraged to optimize general smooth non-convex functions.

7 Discussion

Our results have a number of connections to rates of convergence in classical (smooth) convex optimization and for the power method for symmetric eigenvector computation; here, we explore these in more detail.

7.1 Comparison with convex optimization

For L-smooth α-strongly convex functions, gradient descent finds an ε-suboptimal point within

$$O(1)\cdot\min\left\{\frac{L}{\alpha}\log\frac{LD^2}{\varepsilon}, \frac{LD^2}{\varepsilon}\right\}$$

iterations [17], where D is any constant $D \geq \|x_0 - x^\star\|$, where $x^\star$ a global minimizer. For our (possibly non-convex) problem (1), Corollary 3.3 guarantees that gradient descent finds an ε-suboptimal point (with probability at least $1-\delta$) within

$$O(1) \cdot \min\left\{\frac{L_s}{\rho\|s\| - \gamma}, \frac{L_s\|s\|^2}{\varepsilon}\right\}\left[\log\frac{L_s\|s\|^2}{\varepsilon} + \log\left(1 + \mathbb{I}_{\{\gamma>0\}}\frac{d}{\delta}\right)\right]$$

iterations, where $L_s = \beta + 2\rho\|s\|$. The parallels are immediate: by Corollary 2.5, L_s and $\|s\|$ are precise analogues of L and D in the convex setting. Moreover, the quantity $\rho\|s\| - \gamma$ plays the role of the strong convexity parameter α, but it is well-defined even when f is not convex. When $\lambda^{(1)}(A) = -\gamma \geq 0$, f is $-\gamma$-strongly convex, and because $\rho\|s\| - \gamma > -\gamma$, our analysis for the cubic problem (1) guarantees better conditioning than the generic convex result. The difference between $\rho\|s\| - \gamma$ and $-\gamma$ becomes significant when b is sufficiently large, as we observe from the bounds (6b) and (7). Even in the non-convex case that $\gamma > 0$, gradient descent still exhibits linear convergence whenever high accuracy is desired, that is, when $\varepsilon/\|s\|^2 \leq \rho\|s\| - \gamma$.

When $\gamma > 0$, our guarantee becomes probabilistic and contains a $\log(d/\delta)$ term. Such a term does not appear in results on convex optimization [17], and we believe it is fundamentally related to the non-convexity in the problem: it represents the time necessary to escape from saddle points by means of a random perturbation.

7.2 Comparison with the power method

The power method for finding the bottom eigenvector of A is the recursion $x_{t+1} = (I - (1/\beta)A)x_t / \|(I - (1/\beta)A)x_t\|$ where x_0 is uniform on the unit sphere [13, 16]. This method guarantees that with probability at least $1-\delta$, $x_t^T A x_t \leq -\gamma + \varepsilon$ for all

$$t \geq O(1) \cdot \min\left\{\frac{\beta}{\varepsilon}\log\left(\frac{d}{\delta}\right), \frac{\beta}{\mathsf{gap}}\log\left(\frac{\beta}{\varepsilon}\cdot\frac{d}{\delta}\right)\right\}.$$

When $b = 0$ and $\lambda^{(1)}(A) = -\gamma < 0$, any global minimizer of problem (1) is an eigenvector of A with eigenvalue $-\gamma$ and $\rho\|s\| = \gamma$, so it is natural to compare gradient descent and the power method. For simplicity, let us assume that $\rho = \gamma$ so that $\|s\| = 1$, and both methods converge to unit eigenvectors. Under these assumptions $f(x) = \frac{1}{2}x^T A x + \frac{\gamma}{3}\|x\|^3$ and $f(s) = -\gamma/6$, so $f(x) \leq f(s) + \varepsilon'$ implies

$$\frac{x^T A x}{\|x\|^2} \leq -\frac{\gamma}{3}\left[\frac{1}{\|x\|^2} + 2\|x\|\right] + \frac{2\varepsilon'}{\|x\|^2} \leq -\gamma + \frac{2\varepsilon'}{\|x\|^2}.$$

Consider gradient descent applied to f with a random perturbation as described in Theorem 3.2, with $\overline{\sigma} = 1$. Inspecting the proofs of our theorems (Sec. 4), we see that Lemmas 4.3 and 4.6 imply that with probability at least $1-\delta$ we have $\|\tilde{x}_t\| \geq 1/2$ for every $t \geq O(1)\log(\frac{L_s\|s\|^2}{\varepsilon}\cdot\frac{d}{\delta})$. As in Corollary 3.3, setting $\eta = \frac{1}{4(\beta+\rho R)} = \frac{1}{8\beta}$ guarantees that with probability at least $1-\delta$, $\tilde{x}_t^T A\tilde{x}_t/\|\tilde{x}_t\|^2 \leq -\gamma + \varepsilon$ for all

$$t \geq O(1) \cdot \min\left\{\frac{\beta}{\varepsilon}\log\left(\frac{\beta}{\varepsilon}\cdot\frac{d}{\delta}\right), \frac{\beta}{\sqrt{\varepsilon(\mathsf{gap}\wedge\gamma)}}\log\left(\frac{\beta}{\varepsilon}\cdot\frac{d}{\delta}\right)\right\}.$$

Comparing the rates of convergence, we see that both exhibit the $\log(d/\delta)$ hallmark of non-convexity and gap-free and gap-dependent convergence regimes. Of course, the power method also

finds eigenvectors when $\gamma < 0$, while the unique solution to problem (1) when $b = 0$ and $\gamma < 0$ is simply $s = 0$. In the gap-dependent regime, however, the power method enjoys linear convergence when $\varepsilon < \mathsf{gap}$, while our bounds have a $1/\sqrt{\varepsilon}$ factor. Although this may be due to looseness in our analysis, we suspect it is real and related to the fact that gradient descent needs to "grow" the iterates to have norm $\|x_t\| \approx \gamma_+/\rho$, while the power method iterates always have unit norm. If one is only interested in finding eigenvectors of A, there is probably no reason to prefer the cubic-regularized objective to the power method.

Acknowledgment

YC and JCD were partially supported by the SAIL-Toyota Center for AI Research. YC was partially supported by the Stanford Graduate Fellowship and the Numerical Technologies Fellowship. JCD was partially supported by the National Science Foundation award NSF-CAREER-1553086.

Appendix

For convenience of the reader, throughout the appendix we restate lemmas that appeared in the main text prior to giving their proofs.

A Proof of Lemma 2.2

Before proving Lemma 2.2, we state and prove two technical lemmas (see Sec. A.1 for the proof conditional on these lemmas). For the first lemma, let $\kappa \in \mathbb{R}^d$ satisfy $\kappa^{(1)} \le \kappa^{(2)} \le \ldots \le \kappa^{(d)}$, let ν_t be a nonnegative and nondecreasing sequence, $0 \le \nu_1 \le \nu_2 \le \ldots$, and consider the process

$$z_t^{(i)} = (1 - \kappa^{(i)} - \nu_{t-1}) z_{t-1}^{(i)} + 1. \tag{26}$$

Additionally, assume $1 - \kappa^{(i)} - \nu_{t-1} \ge 0$ for all i and t.

Lemma A.1. *Let $z_0^{(i)} = c_0 \ge 0$ for every $i \in [d]$. Then for every $t \in \mathbb{N}$ and $j \in [d]$, the following holds.*

(i) If $z_t^{(j)} \le z_{t-1}^{(j)}$ then also $z_{t'}^{(j)} \le z_{t'-1}^{(j)}$ for every $t' > t$.

(ii) If $z_t^{(j)} \ge z_{t-1}^{(j)}$, then $z_t^{(j)}/z_{t+1}^{(j)} \ge z_t^{(i)}/z_{t+1}^{(i)}$ for every $i \le j$.

(iii) If $z_{t+1}^{(i)} \le z_t^{(i)}$, then $z_{t+1}^{(j)} \le z_t^{(j)}$ for every $j \ge i$.

Proof. For shorthand, we define $\delta_{t-1}^{(i)} \triangleq \kappa^{(i)} + \nu_{t-1}$.

We first establish part (i) of the lemma. By (26), we have

$$z_{t+1}^{(j)} - z_t^{(j)} = (1 - \delta_{t-1}^{(j)})(z_t^{(j)} - z_{t-1}^{(j)}) - (\delta_t^{(j)} - \delta_{t-1}^{(j)}) z_t^{(j)}$$

By our assumptions that $z_0^{(j)} \ge 0$ and that $1 - \delta_t^{(j)} \ge 0$ for every t we immediately have that $z_t^{(j)} \ge 0$, and therefore also $(\delta_t^{(j)} - \delta_{t-1}^{(j)}) z_t^{(j)} = (\nu_t - \nu_{t-1}) z_t^{(j)} \ge 0$. We therefore conclude that

$$z_{t+1}^{(j)} - z_t^{(j)} \le (1 - \delta_{t-1}^{(j)})(z_t^{(j)} - z_{t-1}^{(j)}) \le 0,$$

and induction gives part (i).

To establish part (ii) of the lemma, first note that by the contrapositive of part (i), $z_t^{(j)} \geq z_{t-1}^{(j)}$ for some t implies $z_{t'}^{(j)} \geq z_{t'-1}^{(j)}$ for any $t' \leq t$. We prove by induction that

$$z_{t'}^{(i)} - z_{t'}^{(j)} \leq (\kappa^{(j)} - \kappa^{(i)}) z_{t'}^{(i)} z_{t'}^{(j)} \tag{27}$$

for any $i \leq j$ and $t' \leq t$. The basis of the induction is immediate from the assumption $z_0^{(i)} = z_0^{(j)} \geq 0$. Assuming the property holds through time $t'-1$ for $t' \leq t$, we obtain

$$\begin{aligned}
\frac{z_{t'}^{(i)} - z_{t'}^{(j)}}{z_{t'}^{(i)} z_{t'}^{(j)}} &= \frac{(1 - \delta_{t'-1}^{(i)})(z_{t'-1}^{(i)} - z_{t'-1}^{(j)}) + (\delta_{t'-1}^{(j)} - \delta_{t'-1}^{(i)}) z_{t'-1}^{(j)}}{z_{t'}^{(i)} z_{t'}^{(j)}} \\
&\leq \frac{(1 - \delta_{t'-1}^{(i)})(\kappa^{(j)} - \kappa^{(i)}) z_{t'-1}^{(i)} z_{t'-1}^{(j)}}{z_{t'}^{(i)} z_{t'}^{(j)}} = (\kappa^{(j)} - \kappa^{(i)}) \frac{z_{t'-1}^{(j)}}{z_{t'}^{(j)}} \leq \kappa^{(j)} - \kappa^{(i)}
\end{aligned}$$

where the first inequality uses inequality (27) (assumed by induction) and the second uses $z_{t'-1}^{(j)} \leq z_{t'}^{(j)}$ for any $t' \leq t$, as argued above. With the bound $z_t^{(i)} - z_t^{(j)} \leq (\kappa^{(j)} - \kappa^{(i)}) z_t^{(i)} z_t^{(j)}$ in place, we may finish the proof of part (ii) by noting that

$$\frac{z_t^{(j)}}{z_{t+1}^{(j)}} - \frac{z_t^{(i)}}{z_{t+1}^{(i)}} = \frac{z_{t+1}^{(i)} z_t^{(j)} - z_{t+1}^{(j)} z_t^{(i)}}{z_{t+1}^{(j)} z_{t+1}^{(i)}} = \frac{(\kappa^{(j)} - \kappa^{(i)}) z_t^{(i)} z_t^{(j)} - (z_t^{(i)} - z_t^{(j)})}{z_{t+1}^{(j)} z_{t+1}^{(i)}} \geq 0.$$

Lastly, we prove part (iii). If $z_t^{(j)} \leq z_{t-1}^{(j)}$ then we have $z_{t+1}^{(j)} \leq z_t^{(j)}$ by part (i). Otherwise we have $z_t^{(j)} \geq z_{t-1}^{(j)}$, and so $z_t^{(j)} / z_{t+1}^{(j)} \geq z_t^{(i)} / z_{t+1}^{(i)}$ by part (ii). As $z_{t+1}^{(i)} \leq z_t^{(i)}$, this implies $z_t^{(j)} / z_{t+1}^{(j)} \geq z_t^{(i)} / z_{t+1}^{(i)} \geq 1$ and therefore $z_{t+1}^{(j)} \leq z_t^{(j)}$ as required. □

Our second technical lemma provides a lower bound on certain inner products in the gradient iteration. In the lemma, we recall the definition (6a) of R.

Lemma A.2. *Assume that $\|x_\tau\|$ is non-decreasing in τ for $\tau \leq t$, that $\|x_t\| \leq R$, and that $x_t^T \nabla f(x_t) \leq 0$. Then $x_t^T A \nabla f(x_t) \geq \beta x_t^T \nabla f(x_t)$.*

Proof. If we define $z_t^{(i)} = x_t^{(i)} / (-\eta b^{(i)})$, then evidently

$$z_{t+1}^{(i)} = (1 - \underbrace{\eta \lambda^{(i)}(A)}_{\triangleq \kappa^{(i)}} - \underbrace{\eta \rho \|x_t\|}_{\triangleq \nu_t}) z_t^{(i)} + 1.$$

We verify that $z_t^{(i)}$ satisfies the conditions of Lemma A.1:

i. By definition $\kappa^{(i)}$ are increasing in i, and $\nu_0 \leq \nu_1 \leq \cdots \leq \nu_t$ by our assumption that $\|x_\tau\|$ is non-decreasing for $\tau \leq t$.

ii. As $\eta \leq 1/(\beta + \rho R)$ for $\tau \leq t$, we have that $\kappa^{(i)} + \nu_\tau \leq 1$ for $\tau \leq t$ and $i \in [d]$.

iii. As $x_0 = -rb/\|b\|$, $z_0^{(i)} = r/(\eta \|b\|) \geq 0$ for every i.

We may therefore apply Lemma A.1, part (iii) to conclude that $z_t^{(i)} - z_{t+1}^{(i)} \geq 0$ implies $z_t^{(j)} - z_{t+1}^{(j)} \geq 0$ for every $j \geq i$. Since $z_t^{(i)} \geq 0$ for every i,

$$\operatorname{sign}\left(x_t^{(i)}\left(x_t^{(i)} - x_{t+1}^{(i)}\right)\right) = \operatorname{sign}\left(z_t^{(i)}\left(z_t^{(i)} - z_{t+1}^{(i)}\right)\right) = \operatorname{sign}\left(z_t^{(i)} - z_{t+1}^{(i)}\right),$$

and there must thus exist some $i^* \in [d]$ such that $x_t^{(i)}(x_t^{(i)} - x_{t+1}^{(i)}) \leq 0$ for every $i \leq i^*$ and $x_t^{(i)}(x_t^{(i)} - x_{t+1}^{(i)}) \geq 0$ for every $i > i^*$. We thus have (by expanding in the eigenbasis of A) that

$$\begin{aligned}
x_t^T A \nabla f(x_t) &= \frac{1}{\eta}\sum_{i=1}^{i^*} \lambda^{(i)}(A)\, x_t^{(i)}\left(x_t^{(i)} - x_{t+1}^{(i)}\right) + \frac{1}{\eta}\sum_{i=i^*+1}^{d} \lambda^{(i)}(A)\, x_t^{(i)}\left(x_t^{(i)} - x_{t+1}^{(i)}\right) \\
&\geq \lambda^{(i^*)}(A)\frac{1}{\eta}\sum_{i=1}^{i^*} x_t^{(i)}\left(x_t^{(i)} - x_{t+1}^{(i)}\right) + \lambda^{(i^*+1)}(A)\frac{1}{\eta}\sum_{i=i^*+1}^{d} x_t^{(i)}\left(x_t^{(i)} - x_{t+1}^{(i)}\right) \\
&\geq \lambda^{(i^*)}(A)\frac{1}{\eta}\sum_{i=1}^{d} x_t^{(i)}\left(x_t^{(i)} - x_{t+1}^{(i)}\right) = \lambda^{(i^*)}(A)\, x_t^T \nabla f(x_t) \geq \beta x_t^T \nabla f(x_t)
\end{aligned}$$

where the first two inequalities use the fact the $\lambda^{(i)}$ is non-decreasing with i, and the last inequality uses our assumption that $x_t^T \nabla f(x_t) \leq 0$ along with $\lambda^{(d)}(A) \leq \beta$. □

A.1 Proof of Lemma 2.2

Lemma 2.2. *Let Assumptions A and B hold. Then the iterates* (8) *of gradient descent satisfy* $x_t^T \nabla f(x_t) \leq 0$*, the norms* $\|x_t\|$ *are non-decreasing, and* $\|x_t\| \leq R$*.*

By definition of the gradient iteration, we have

$$\|x_{t+1}\|^2 = \|x_t\|^2 - 2\eta x_t^T \nabla f(x_t) + \eta^2 \|\nabla f(x_t)\|^2, \tag{28}$$

and therefore if we can show that $x_t^T \nabla f(x_t) \leq 0$ for all t, the lemma holds. We give a proof by induction. The basis of the induction $x_0^T \nabla f(x_0) \leq 0$ is immediate as $r \mapsto f(-rb/\|b\|)$ is decreasing until $r = R_c$ (recall the definition (7)), and $x_0^T \nabla f(x_0) = 0$ for $r \in \{0, R_c\}$. Our induction assumption is that $x_{t'-1}^T \nabla f(x_{t'-1}) \leq 0$ (and hence also $\|x_{t'}\| \geq \|x_{t'-1}\|$) for $t' \leq t$ and we wish to show that $x_t^T \nabla f(x_t) \leq 0$. Note that

$$x^T \nabla f(x) = x^T A x + \rho \|x\|^3 + b^T x \geq \rho \|x\|^3 - \gamma \|x\|^2 - \|b\| \|x\|$$

and therefore $x^T \nabla f(x) > 0$ for every $\|x\| > R_{\text{low}} \triangleq \frac{\gamma}{2\rho} + \sqrt{\left(\frac{\gamma}{2\rho}\right)^2 + \frac{\|b\|}{\rho}}$. Therefore, our induction assumption also implies $\|x_{t'-1}\| \leq R_{\text{low}} \leq R$ for every $t' \leq t$.

Using that $\nabla^2 f$ is 2ρ-Lipschitz, a Taylor expansion immediately implies [18, Lemma 1] that for all vectors Δ, we have

$$\left\|\nabla f(x+\Delta) - (\nabla f(x) + \nabla^2 f(x)\Delta)\right\| \leq \rho \|\Delta\|^2. \tag{29}$$

Thus, if we define $\Delta_t \triangleq \frac{1}{\eta^2}[\nabla f(x_t) - (\nabla f(x_{t-1}) - \eta \nabla^2 f(x_{t-1}) \nabla f(x_{t-1}))]$, we have $\|\Delta_t\| \le \rho \|\nabla f(x_{t-1})\|^2$, and using the iteration $x_t = x_{t-1} - \eta \nabla f(x_{t-1})$ yields

$$x_t^T \nabla f(x_t) = x_{t-1}^T \nabla f(x_{t-1}) - \eta \|\nabla f(x_{t-1})\|^2 - \eta \underbrace{x_{t-1}^T \nabla^2 f(x_{t-1}) \nabla f(x_{t-1})}_{\triangleq \mathcal{T}_1} + \eta^2 \underbrace{\nabla f(x_{t-1}) \nabla^2 f(x_{t-1}) \nabla f(x_{t-1})}_{\triangleq \mathcal{T}_2} + \eta^2 \underbrace{x_t^T \Delta_t}_{\triangleq \mathcal{T}_3}. \tag{30}$$

We bound each of the terms $\mathcal{T}_i$ in turn. We have that

$$\begin{aligned}\mathcal{T}_1 = x_{t-1}^T \nabla^2 f(x_{t-1}) \nabla f(x_{t-1}) &= x_{t-1}^T A \nabla f(x_{t-1}) + 2\rho \|x_{t-1}\| x_{t-1}^T \nabla f(x_{t-1}) \\ &\ge (\beta + 2\rho \|x_{t-1}\|) x_{t-1}^T \nabla f(x_{t-1}) \ge (\beta + 2\rho R) x_{t-1}^T \nabla f(x_{t-1}),\end{aligned}$$

where both inequalities follow from induction assumption; the first is Lemma A.2 and the second is due to $\|x_{t-1}\| \le R$ and $x_{t-1}^T \nabla f(x_{t-1}) \le 0$.

Treating the second order term $\mathcal{T}_2$, we obtain that

$$\mathcal{T}_2 \le \left\|\nabla^2 f(x_{t-1})\right\|_{\text{op}} \|\nabla f(x_{t-1})\|^2 \le (\beta + 2\rho R) \|\nabla f(x_{t-1})\|^2,$$

and, by the Lipschitz bound (29), the remainder term $\mathcal{T}_3$ satisfies

$$\begin{aligned}\mathcal{T}_3 = x_t^T \Delta_t \le \|x_t\| \|r\| &\le \rho \|x_t\| \|\nabla f(x_{t-1})\|^2 \le \rho \|x_{t-1} - \eta \nabla f(x_{t-1})\| \|\nabla f(x_{t-1})\|^2 \\ &\le \rho \|x_{t-1}\| \|\nabla f(x_{t-1})\|^2 + \rho\eta \|\nabla f(x_{t-1})\|^3.\end{aligned}$$

Using that $\|\nabla f(x)\| = \|\nabla f(x) - \nabla f(s)\| \le (\beta + 2R) \|x - s\| \le R(\beta + 2\rho R)$ for $\|x\| \le R$ and that $\eta \le 1/2(\beta + 2\rho R)$, our inductive assumption that $\|x_{t-1}\| \le R$ thus guarantees that $\mathcal{T}_3 \le 2\rho R \|\nabla f(x_{t-1})\|^2$. Combining our bounds on the terms $\mathcal{T}_i$ in expression (30), we have that

$$x_t^T \nabla f(x_t) \le (1 - \eta(\beta + 2\rho R)) x_{t-1}^T \nabla f(x_{t-1}) - \left(\eta - \eta^2(\beta + 4\rho R)\right) \|\nabla f(x_{t-1})\|^2.$$

Using $\eta \le 1/(\beta + 4\rho R)$ shows that $x_t^T \nabla f(x_t) \le 0$, completing our induction. By the expansion (28), we have $\|x_t\| \le \|x_{t+1}\|$ as desired, and that $x_t^T \nabla f(x_t) \le 0$ for all t guarantees that $\|x_t\| \le R_{\text{low}} \le R$.

B Proofs of technical results from Section 4

As in the statement of our major theorems and as we note in the beginning of Section 4, we tacitly assume Assumptions A and B throughout this section.

B.1 Proof of Lemma 4.1

Lemma 4.1. *For each $t > 0$, we have*

$$\|x_t - s\|^2 \le \left(1 - \eta \left[\rho \|x_t\| - \left(\gamma - \frac{\rho \|s\| - \gamma}{2}\right)\right]\right) \|x_{t-1} - s\|^2$$

Expanding $x_t = x_{t-1} - \eta \nabla f(x_{t-1})$, we have

$$\|x_t - s\|^2 = \|x_{t-1} - s\|^2 - 2\eta (x_{t-1} - s)^T \nabla f(x_{t-1}) + \eta^2 \|\nabla f(x_{t-1})\|^2. \tag{31}$$

Using the equality $\nabla f(x) = A_s(x-s)+\rho(\|s\|-\|x\|)x$, we rewrite the cross-term $(x_{t-1}-s)^T\nabla f(x_{t-1})$ as

$$\begin{aligned}(x_{t-1}-s)^T A_s(x_{t-1}-s) &+ \rho(\|x_{t-1}\|-\|s\|)(\|x_{t-1}\|^2 - s^T x_{t-1}) \\ &= (x_{t-1}-s)^T\left(A_s + \frac{\rho}{2}(\|x_{t-1}\|-\|s\|)I\right)(x_{t-1}-s) \\ &\quad + \frac{\rho}{2}(\|s\|-\|x_{t-1}\|)^2(\|x_{t-1}\|+\|s\|). \end{aligned} \tag{32}$$

Moving to the second order term $\|\nabla f(x_{t-1})\|^2$ from the expansion (31), we find

$$\begin{aligned}\|\nabla f(x_{t-1})\|^2 &= \|A_s(x_{t-1}-s)+\rho(\|x_{t-1}\|-\|s\|)x_{t-1}\|^2 \\ &\le 2(x_{t-1}-s)^T A_s^2(x_{t-1}-s) + 2\rho^2(\|x_{t-1}\|-\|s\|)^2\|x_{t-1}\|^2.\end{aligned}$$

Combining this inequality with the cross-term calculation (32) and the squared distance (31) we obtain

$$\begin{aligned}\|x_t-s\|^2 \le (x_{t-1}-s)^T(I-2\eta A_s(I-\eta A_s)-\eta\rho(\|x_{t-1}\|-\|s\|)I)(x_{t-1}-s) \\ -\eta\rho(\|s\|-\|x_{t-1}\|)^2(\|x_{t-1}\|(1-2\eta\rho\|x_{t-1}\|)+\|s\|).\end{aligned}$$

Using $\eta \le \frac{1}{4(\beta+\rho R)} \le \frac{1}{4}\|A_s\|_{\mathrm{op}}$ yields $2\eta A_s(1-\eta A_s) \succeq \frac{3}{2}\eta A_s \succeq \frac{3}{2}(-\gamma+\rho\|s\|)I$, so

$$\|x_t-s\|^2 \le \left(1-\frac{\eta}{2}[-3\gamma+\rho(\|s\|+2\|x_{t-1}\|)]\right)\|x_{t-1}-s\|^2 - \eta\rho(\|s\|-\|x_{t-1}\|)^2\|s\|.$$

B.2 Proof of Lemma 4.3

Lemma 4.3. *Let $\delta > 0$. Then $\rho\|x_t\| \ge \gamma-\delta$ for all $t \ge \frac{2}{\eta\delta}\log(1+\frac{\gamma_+^2}{4\rho|b^{(1)}|})$.*

The claim is trivial when $\gamma \le 0$ as it clearly implies $\rho\|x_t\| \ge \gamma$, so we assume $\gamma_+ = \gamma > 0$. Using Proposition 2.4 that gradient descent is convergent, we may define $t^* = \min\{t : \rho\|x_t\| \le \gamma-\delta\}$. Then for every $t \le t^*$, the gradient iteration (8) satisfies

$$\begin{aligned}\frac{x_t^{(1)}}{-\eta b^{(1)}} &= (1+\eta\gamma-\eta\rho\|x_{t-1}\|)\frac{x_{t-1}^{(1)}}{-\eta b^{(1)}} + 1 \\ &\ge (1+\eta\delta)\frac{x_{t-1}^{(1)}}{-\eta b^{(1)}} + 1 \ge \cdots \ge \frac{1}{\eta\delta}\left((1+\eta\delta)^t - 1\right).\end{aligned}$$

Multiplying both sides of the equality by $\eta|b^{(1)}|$ and using that $x_t^{(1)}b^{(1)} \le 0$, we have

$$\frac{\gamma-\delta}{\rho} \ge \|x_{t^*}\| \ge |x_{t^*}^{(1)}| \ge \frac{|b^{(1)}|}{\delta}\left((1+\eta\delta)^{t^*}-1\right).$$

Consequently,

$$t^* \le \frac{\log\left(1+\frac{(\gamma-\delta)\delta}{\rho|b^{(1)}|}\right)}{\log(1+\eta\delta)} \le \frac{2}{\eta\delta}\log\left(1+\frac{\gamma_+^2}{4\rho|b^{(1)}|}\right),$$

where we used $\eta\delta \le \eta\gamma \le \gamma/\beta \le 1$, whence $\log(1+\eta\delta) \ge \frac{\eta\delta}{2}$, and $\gamma\delta-\delta^2 \le \frac{\gamma_+^2}{4}$.

B.3 Proof of Lemma 4.4

Lemma 4.4. *Let Π be any projection matrix satisfying $\Pi A = A\Pi$ for which $\Pi A_s \succeq \delta\Pi$ for some $\delta > 0$. For all $t > 0$,*

$$\begin{aligned}\left\|\Pi A_s^{1/2}(x_t - s)\right\|^2 &\leq (1-\eta\delta)\left\|\Pi A_s^{1/2}(x_{t-1}-s)\right\|^2 \\ &\quad + \sqrt{8}\eta\rho\left(\|s\| - \|x_{t-1}\|\right)\left[\rho\left(\|s\| - \|x_{t-1}\|\right)\|x_{t-1}\|^2 + \left\|(I-\Pi)A_s\right\|_{\mathrm{op}}\|s\|^2\right].\end{aligned}$$

For typographical convenience, we prove the result with $t+1$ replacing t. Using the commutativity of Π and A, we have $\Pi A_s = A_s\Pi$, so

$$\begin{aligned}\left\|\Pi A_s^{1/2}\left(x_{t+1} - s\right)\right\|^2 &= \left\|\Pi A_s^{1/2}\left(x_t - s\right)\right\|^2 \\ &\quad - 2\eta\left(x_t - s\right)^T A_s\Pi\nabla f\left(x_t\right) + \eta^2\left\|\Pi A_s^{1/2}\nabla f\left(x_t\right)\right\|^2. \end{aligned} \tag{33}$$

We substitute $\nabla f(x) = A_s(x-s) - \rho\left(\|s\| - \|x\|\right)x$ in the cross term to obtain

$$\left(x_t - s\right)^T\Pi A_s\nabla f\left(x_t\right) = \left(x_t - s\right)^T\Pi A_s^2\Pi\left(x_t - s\right) - \rho\left(\|s\| - \|x_t\|\right)x_t^T\Pi A_s\left(x_t - s\right).$$

Substituting $A_s(x-s) = \nabla f(x) + \rho\left(\|s\| - \|x\|\right)x$ in the last term yields

$$x_t^T\Pi A_s\left(x_t - s\right) = x_t^T\Pi\nabla f\left(x_t\right) + \rho\left(\|s\| - \|x_t\|\right)\|\Pi x_t\|^2. \tag{34}$$

Invoking Corollary 2.5 and the fact that $x_t^T\nabla f(x_t) \leq 0$, we get

$$\begin{aligned}x_t^T\Pi\nabla f\left(x_t\right) &= x_t^T\nabla f\left(x_t\right) - x_t^T\left(I-\Pi\right)\nabla f\left(x_t\right) \\ &\leq -x_t^T\left(I-\Pi\right)\nabla f\left(x_t\right) \\ &= -x_t^T\left(I-\Pi\right)A_s\left(x_t - s\right) + \rho\left(\|s\| - \|x_t\|\right)\left\|\left(I-\Pi\right)x_t\right\|^2 \\ &\leq \left\|\left(I-\Pi\right)A_s\right\|_{\mathrm{op}}\|x_t\|\,\|x_t - s\| + \rho\left(\|s\| - \|x_t\|\right)\left\|\left(I-\Pi\right)x_t\right\|^2 \\ &\leq \sqrt{2}\left\|\left(I-\Pi\right)A_s\right\|_{\mathrm{op}}\|s\|^2 + \rho\left(\|s\| - \|x_t\|\right)\left\|\left(I-\Pi\right)x_t\right\|^2,\end{aligned}$$

where in the last line we used $x_t^T s \geq 0$ (by Lemma 2.3). Combining this with the cross terms (34), we find that

$$x_t^T\Pi A_s\left(x_t - s\right) \leq \sqrt{2}\left\|\left(I-\Pi\right)A_s\right\|_{\mathrm{op}}\|s\|^2 + \rho\left(\|s\| - \|x_t\|\right)\|x_t\|^2.$$

Moving on to the second order term in the expansion (33), we have

$$\begin{aligned}\left\|\Pi A_s^{1/2}\nabla f\left(x_t\right)\right\|^2 &= \left\|\Pi A_s^{3/2}\left(x_t - s\right) + \rho\left(\|x_t\| - \|s\|\right)A_s^{1/2}\Pi x_t\right\|^2 \\ &\leq 2\left\|\Pi A_s^{3/2}\left(x_t - s\right)\right\|^2 + 2\rho^2\left\|\Pi A_s\right\|_{\mathrm{op}}\left(\|x_t\| - \|s\|\right)^2\|x_t\|^2.\end{aligned}$$

Combining the previous two displays into the expansion (33), we have

$$\begin{aligned}\left\|\Pi A_s^{1/2}\left(x_{t+1} - s\right)\right\|^2 &\leq \left(x_t - s\right)^T\left(I - 2\eta\Pi A_s\left(I - \eta\Pi A_s\right)\right)\Pi A_s\left(x_t - s\right) \\ &\quad + 2\eta\rho\left(\|s\| - \|x_t\|\right)\left[\sqrt{2}\left\|\left(I-\Pi\right)A_s\right\|_{\mathrm{op}}\|s\|^2\right. \\ &\quad \left. + \left(1 + \eta\left\|\Pi A_s\right\|_{\mathrm{op}}\right)\rho\left(\|x_t\| - \|s\|\right)\|x_t\|^2\right].\end{aligned}$$

Using $\eta \le 1/4\,(\beta + \rho R)$ which guarantees $0 \preceq \eta\Pi A_s \preceq (\sqrt{2}-1)I \prec I/2$ together with the assumption that $\Pi A_s \succeq \delta\Pi$ gives

$$0 \preceq I - 2\eta\Pi A_s\,(I - \eta\Pi A_s) \preceq (1-\eta\delta)I$$

and therefore

$$\begin{aligned}\left\|\Pi A_s^{1/2}\,(x_{t+1}-s)\right\|^2 &\le (1-\eta\delta)\left\|\Pi A_s^{1/2}\,(x_t - s)\right\|^2 \\ &\quad + \sqrt{8}\eta\rho\,(\|s\| - \|x_t\|)\left[\rho\,(\|s\| - \|x_t\|)\,\|x_t\|^2 + \|(I-\Pi)\,A_s\|_{\mathrm{op}}\,\|s\|^2\right].\end{aligned}$$

B.4 Proof of Lemma 4.5

Lemma 4.5. *Let $t \ge 0$, $\delta \ge 0$, and define $\bar{\delta} = \delta \vee \mathsf{gap}$. If $\rho\,\|s\| \le \gamma + \sqrt{\delta\bar{\delta}}$ and $\rho\,\|x_t\| \ge \gamma - \frac{1}{3}\sqrt{\delta\bar{\delta}}$, then for any $\tau \ge 0$,*

$$\begin{aligned}\left\|\Pi_\delta A_s^{1/2}\,(x_{t+\tau}-s)\right\|^2 &\le \left(1-\eta\bar{\delta}\right)^\tau\left\|\Pi_\delta A_s^{1/2}\,(x_t-s)\right\|^2 + 13\,\|s\|^2\,\delta \\ &\le 2\,(\beta+\rho\,\|s\|)\,\|s\|^2\,e^{-\eta\bar{\delta}\tau} + 13\,\|s\|^2\,\delta.\end{aligned}$$

The conditions of the lemma imply that for $\tau \ge 0$,

$$\rho(\|s\| - \|x_{t+\tau}\|) \le 4\sqrt{\delta\bar{\delta}}/3$$

and also that $\|(I-\Pi_\delta)\,A_s\|_{\mathrm{op}} \le 2\sqrt{\delta\bar{\delta}}$ and $\Pi_\delta A_s \succeq \bar{\delta} I$. Substituting these bounds into Lemma 4.4 along with $\|x_{t-1}\| \le \|s\|$ (Corollary 2.5), we get

$$\left\|\Pi_\delta A_s^{1/2}\,(x_{t+\tau}-s)\right\|^2 \le \left(1-\eta\bar{\delta}\right)\left\|\Pi_\delta A_s^{1/2}\,(x_{t+\tau-1}-s)\right\|^2 + 13\eta\delta\bar{\delta}\,\|s\|^2.$$

Iterating this τ times gives

$$\begin{aligned}\left\|\Pi_\delta A_s^{1/2}\,(x_{t+\tau}-s)\right\|^2 &\le \left(1-\eta\bar{\delta}\right)^\tau\left\|\Pi_\delta A_s^{1/2}\,(x_t-s)\right\|^2 + 13\delta\,\|s\|^2\left(1-\left(1-\eta\bar{\delta}\right)^\tau\right) \\ &\le 2\,(\beta+\rho\,\|s\|)\,\|s\|^2\,e^{-\eta\bar{\delta}\tau} + 13\,\|s\|^2\,\delta\end{aligned}$$

where the last transition uses that

$$\left\|\Pi_\delta A_s^{1/2}\,(x_t-s)\right\|^2 \le \|A_s\|_{\mathrm{op}}\,\|x_t-s\|^2 \le (\beta+\rho\,\|s\|)\,2\,\|s\|^2.$$

B.5 Proof of Lemma 4.6

Lemma 4.6. *Set $\tilde{b} = b + \sigma q$, where q is uniform on the unit sphere in $\mathbb{R}^d$ and $\sigma > 0$. Let $\tilde{f}(x) = \frac{1}{2}x^TAx + \tilde{b}^Tx + \frac{1}{3}\rho\,\|x\|^3$ and let $\tilde{s}$ be a global minimizer of $\tilde{f}$. Then*

(i) For $d > 2$, $\mathbb{P}(|\tilde{b}^{(1)}| \le \sqrt{\pi}\sigma\delta/\sqrt{2d}) \le \delta$

(ii) $|f(x) - \tilde{f}(x)| \le \sigma\,\|x\|$ for all $x \in \mathbb{R}^d$

(iii) $\left|\,\|s\|^2 - \|\tilde{s}\|^2\,\right| \le 2\sigma/\rho$.

To establish part (i) of the lemma, note that marginally $[q^{(1)}]^2 \sim \text{Beta}(\frac{1}{2}, \frac{d-1}{2})$ and that $q^{(1)}$ is symmetrically distributed. Therefore, for $d > 2$ the density of $\tilde{b}^{(1)} = b^{(1)} + \sigma q^{(1)}$ is maximal at $b^{(1)}$ and is monotonically decreasing in the distance from $b^{(1)}$. Therefore we have

$$\mathbb{P}\left(|\tilde{b}^{(1)}| \leq \sigma\sqrt{\pi}\delta/\sqrt{2d}\right) \leq \mathbb{P}\left(|q^{(1)}| \leq \sqrt{\pi}\delta/\sqrt{2d}\right) \leq \delta,$$

where the bound $p_1(\gamma) \leq \sqrt{d/(2\pi\gamma)}$ on the density p_1 of $q^{(1)}$ yields the last inequality.

Part (iii) of the lemma is immediate, as

$$|f(x) - \tilde{f}(x)| = |(b - \tilde{b})^T x| \leq \sigma \|q\| \|x\| = \sigma \|x\|.$$

Part (ii) of the lemma follows by viewing $\|s\|^2$ as a function of b and noting that $b \mapsto \|s\|^2$ is $2/\rho$-Lipschitz continuous. To see this claim, we use the inverse function theorem. First, note that $\|s\|^2$ is a well-defined function of b, because s is not unique only when $\|s\| = \gamma/\rho$. Next, from the relation $b = -A_s s$ we see that the inverse mapping $s \mapsto b$ is a smooth function, with Jacobian

$$\frac{\partial b}{\partial s} = -A_s - \rho\frac{ss^T}{\|s\|} = -\nabla^2 f(s).$$

Let us now evaluate $\partial \|s\|^2 / \partial b$ when the mapping $s \mapsto b(s) = -(A + \rho\|s\| I)s$ is invertible (*i.e.* in the case that $\|s\| > \gamma/\rho$); the inverse function theorem yields

$$\frac{\partial \|s\|^2}{\partial b} = \frac{\partial \left(s^T s\right)}{\partial b} = 2\frac{\partial s}{\partial b}s = -2\left(\nabla^2 f(s)\right)^{-1} s.$$

The mapping $s \mapsto (\nabla^2 f(s))^\dagger s$ is continuous in s even when $A_s \succeq 0$ is singular, and therefore the preceding expression is valid (as the natural limit) when $\|s\| \to \gamma/\rho$. Moreover, since $\nabla^2 f(s) \succeq \rho ss^T/\|s\|$, we have

$$\left\|\frac{\partial \|s\|^2}{\partial b}\right\| = 2\left\|\left(\nabla^2 f(s)\right)^{-1} s\right\| \leq 2\left\|\left(\rho ss^T/\|s\|\right)^{-1} s\right\| = \frac{2}{\rho}.$$

We thus conclude that $b \mapsto \|s\|^2$ is a $2/\rho$-Lipschitz continuous function of b, and therefore $|\|s\|^2 - \|\tilde{s}\|^2| \leq (2/\rho)\|b - \tilde{b}\| = 2\sigma/\rho$.

C Proof of Proposition 5.1

We begin with a lemma implicitly assuming the conditions of Proposition 5.1.

Lemma C.1. *For all t we have $\|x_t\| \leq 2R$, with R given by* (6a).

Proof. Note that R minimizes the polynomial $-\|b\|r - \beta r^2/2 + \rho r^3/3$ as it solves $-\|b\| - \beta R + \rho R^2 = 0$. This implies that for every $\|x\| > 2R$ we have

$$f(x) \geq -\|b\|\|x\| - \frac{\beta}{2}\|x\|^2 + \frac{\rho}{3}\|x\|^3 > 2R\left(-\|b\| - \beta R + \frac{4\rho}{3}R^2\right) = \frac{2\rho}{3}R^3 \geq 0,$$

where the first inequality follows because $b^T x \geq -\|b\|\|x\|$ and $\beta \geq \|A\|_{\text{op}}$, the second because $-\|b\|\|x\| - \beta\|x\|^2/2 + \rho\|x\|^3/3$ is increasing in $\|x\|$ for $\|x\| \geq R$, and in the last inequality we substituted $\|b\| = \rho R^2 - \beta R$. By Assumption B, $f(x_0) \leq 0$, and the definition (22) of the step size η_t

guarantees that $f(x_t)$ is non-increasing. Thus $f(x_t) \le 0$ for all t, so $\|x_t\| \le 2R$. □

As in our proof of Lemma 2.3, we focus on the on the first coordinate of the iteration (21) (*i.e.* $x_{t+1} = x_t - \eta_t \nabla f(x_t)$) in the eigenbasis of A, writing

$$x_{t+1}^{(1)} = \left(1 - \eta_t \left[-\gamma + \rho\|x_t\|\right]\right) x_t^{(1)} - \eta_t b^{(1)}.$$

By the constrains in the definition (22) of the step size η_t, we have

$$1 - \eta_t(-\gamma + \rho\|x_t\|) \ge 1 - \eta_t \left[\frac{\nabla f(x_t)^T A \nabla f(x_t)}{\|\nabla f(x_t)\|^2} + \rho\|x_t\|\right]_+ \ge 0.$$

By Assumption B, $b^{(1)} x_0^{(1)} \le 0$, so $b^{(1)} x_t^{(1)} \le 0$ for every t. Since $u^T A u / \|u\|^2 \le \|A\|_{\mathrm{op}} \le \beta$ for all u and $\|x_t\| \le 2R$ for every t, the step size $\eta_{\mathrm{feas}} \triangleq 1/(\beta + 4\rho R)$ is always feasible, and we have $f(x_{t+1}) \le f(x_t - \eta_{\mathrm{feas}} \nabla f(x_t))$. Moreover, since f is $\beta + 4\rho R$-smooth on the set $\mathbb{B}_{2R} = \{x \in \mathbb{R}^d : \|x\| \le 2R\}$, and as $x_t \in \mathbb{B}_{2R}$ for all t by Lemma C.1, we have $f(x_{t+1}) \le f(x_t) - \frac{\eta_{\mathrm{feas}}}{2}\|\nabla f(x_t)\|^2$, which implies $\nabla f(x_t) \to 0$. Having established $b^{(1)} x_t^{(1)} \le 0$ for every t and $\nabla f(x_t) \to 0$ as $t \to \infty$, the remainder of the proof is identical to that of Proposition 2.4.

D Proofs from Section 6

D.1 Proof of Lemma 6.1

Lemma 6.1. *Let $A \in \mathbb{R}^{d\times d}$ satisfy $\|A\|_{\mathrm{op}} \le \beta$, $b \in \mathbb{R}^d$, $\rho > 0, \|\hat{s}\| > 0$, $\varepsilon' \in (0,1), \delta \in (0,1)$ and $\eta \le 1/(8\beta + 4\rho\|\hat{s}\|)$. With probability at least $1-\delta$, if*

$$\|s\| \ge \|\hat{s}\| \text{ or } \|b\| \ge \max\{\sqrt{\beta\rho}\|\hat{s}\|^{3/2}, \rho\|\hat{s}\|^2\}$$

then $x =$ SOLVE-SUBPROBLEM$(A, b, \rho, \eta, \|\hat{s}\|, \varepsilon', \delta)$ satisfies $f(x) \le (1-\varepsilon')\rho\|\hat{s}\|^3/6$.

For x_0 defined in line 2 of Alg. 2, we have $f(x_0) = -(1/2)R_c\|b\| - (\rho/6)R_c^3$, where R_c is the Cauchy radius (7). Therefore a sufficient condition for $f(x_0) \le -(1-\varepsilon')\rho\|\hat{s}\|^3/6$ is $R_c\|b\| \ge \rho\|\hat{s}\|^3/3$. We have

$$R_c \ge \frac{-\beta}{2\rho} + \sqrt{\left(\frac{\beta}{2\rho}\right)^2 + \frac{\|b\|}{\rho}} \ge \left(\frac{2\|b\|}{3\beta} \wedge \sqrt{\frac{\|b\|}{3\rho}}\right) \ge \frac{1}{3}\left(\frac{\|b\|}{\beta} \wedge \sqrt{\frac{\|b\|}{\rho}}\right),$$

where the second inequality follows from $\sqrt{1+\alpha} \ge 1 + (\alpha/3 \wedge \sqrt{\alpha/3})$ for every $\alpha \ge 0$. Thus, Alg. 2 returns x_0 whenever $\|b\|(\frac{\|b\|}{\beta} \wedge \sqrt{\frac{\|b\|}{\rho}}) \ge \rho\|\hat{s}\|^3$, which is equivalent to the second part of the "or" condition in the lemma.

Now, suppose that the algorithm does not return x_0, *i.e.* $f(x_0) > -(1-\varepsilon')\rho\|\hat{s}\|^3/6$. Since $f(x_0) < -(\rho/6)R_c^3$, this implies that $R_c < \|\hat{s}\|$. Since $\rho R_c \ge \rho R - \beta$, with R defined in (6a), we have $\rho\|\hat{s}\| > \rho R - \beta$. Therefore, $\eta \le 1/(8\beta + 4\rho\|\hat{s}\|) \le 1/(4\beta + 4\rho R)$, and so the stepsize η in gradient descent satisfies Assumption A. Since we choose $\tilde{x}_0$ in accordance with Assumption B, we may invoke Theorem 3.2 with $\varepsilon = \rho\|s\|^3\varepsilon'/12$.

Our setting $\sigma = \frac{\rho^2\|\hat{s}\|^3\varepsilon'}{144(\beta+2\rho\|\hat{s}\|)} = \frac{\rho\|\hat{s}\|^3\varepsilon}{12\|s\|^3(\beta+2\rho\|\hat{s}\|)}$ implies

$$\overline{\sigma} = \frac{12\sigma(\beta + 2\rho\|s\|)}{\rho\varepsilon} = \frac{\beta + 2\rho\|s\|}{\beta + 2\rho\|\hat{s}\|} \cdot \frac{\|\hat{s}\|^3}{\|s\|^3}.$$

Therefore, assuming $\|\hat{s}\| \le \|s\|$, we have that $(\|\hat{s}\|/\|s\|)^3 \le \overline{\sigma} \le 1$. Substituting these upper and lower bounds, Theorem 3.2 shows that, with probability at least $1-\delta$, $f(\tilde{x}_t) \le f(s) + (1+\overline{\sigma})\varepsilon \le f(s) + 2\varepsilon$ for all

$$t \ge \frac{20\|s\|^2}{\eta\varepsilon}\left(6\log\left(1+\frac{3\sqrt{d}}{\delta}\cdot\frac{\|s\|^3}{\|\hat{s}\|^3}\right)+14\log\left(\frac{(\beta+2\rho\|s\|)\|s\|^2}{\varepsilon}\right)\right) \triangleq \tilde{T}_\varepsilon^{\mathrm{sub}}.$$

Using $1/\eta \le 2(\beta+2\rho\|s\|)$ and plugging in $\varepsilon = \rho\|s\|^3\varepsilon'/12$, we see that

$$\begin{aligned}\tilde{T}_\varepsilon^{\mathrm{sub}} &\le \frac{240}{\eta\|s\|\varepsilon'}\left(6\log\left(1+\frac{\sqrt{d}}{\delta}\right)+6\log\left(\left[\frac{1}{\eta\rho\|\hat{s}\|}\right]^3\right)+14\log\left(\frac{6}{\eta\rho\|s\|\varepsilon'}\right)\right)\\ &\le \frac{240}{\eta\|\hat{s}\|\varepsilon'}\left(6\log\left(1+\frac{\sqrt{d}}{\delta}\right)+32\log\left(\frac{6}{\eta\rho\|\hat{s}\|\varepsilon'}\right)\right),\end{aligned}$$

where we used $\|\hat{s}\| \le \|s\|$ and $\varepsilon' < 1$. Therefore T defined in line 4 is larger than $\tilde{T}_\varepsilon^{\mathrm{sub}}$, so with probability at least $1-\delta$, there exists $t \le T$ for which $f(\tilde{x}_t) \le f(s) + \rho\|s\|^3\varepsilon'/6$. Recalling that $f(s) \le -\rho\|s\|^3/6 \le -\rho\|\hat{s}\|^3/6$ by the bound (4a) completes the proof.

D.2 Proof of Proposition 6.2

Proposition 6.2. *Let g satisfy Assumption C, $y_0 \in \mathbb{R}^d$ be arbitrary, and let $\delta \in (0,1]$ and $\epsilon \le \min\{\beta^2/\rho, \rho^{1/3}\}$. With probability at least $1-\delta$, Algorithm 1 finds an ϵ-second-order stationary point (24) in at most*

$$O(1)\cdot\frac{\beta\lceil g(y_0)-g^\star\rceil}{\epsilon^2}\log\left(\frac{d}{\delta}\cdot\frac{\beta\lceil g(y_0)-g^\star\rceil}{\epsilon^2}\right) \tag{25}$$

gradient and Hessian-vector product evaluations.

We always call Solve-subproblem with $\varepsilon' = 1/2$ and $\|\hat{s}\| = \sqrt{\epsilon/(9\rho)}$. As $\epsilon \le \beta^2/\rho$ we have that $\eta = 1/(10\beta) \le 1/(8\beta + 4\rho\|\hat{s}\|)$. Since $\left\|\nabla^2 g(x)\right\|_{\mathrm{op}} \le \beta$ by Assumption C, we conclude that Lemma 6.1 applies to each call of Solve-subproblem. Because each iteration of Alg. 1 guarantees progress $K_{\mathrm{prog}}\epsilon^{3/2}\rho^{-1/2}$ (recall Line 6 of Alg. 1), the algorithm runs for at most $K_{\max} = \lceil 324(g(y_0)-g^\star)\sqrt{\rho}\epsilon^{-3/2}\rceil$ iterations. Letting $\mathcal{E}$ be the event that at each call to Solve-subproblem, the conclusions of Lemma 6.1 hold, the definition (Line 3) of δ' and a union bound guarantee $\mathbb{P}(\mathcal{E}) \ge 1-\delta$. We perform our subsequent analysis deterministically conditional on the event $\mathcal{E}$.

Let f_k be the cubic-regularized quadratic model at iteration k. We call the iteration successful (Line 6) whenever Solve-subproblem finds a point Δ_k such that

$$f(\Delta_k) \le -(1-\varepsilon')\rho\|\hat{s}\|^3/6 = -\frac{1}{2}\left(\frac{\epsilon}{9\rho}\right)^{3/2}\frac{\rho}{6} = -K_{\mathrm{prog}}\epsilon^{3/2}\rho^{-1/2}.$$

The bound (23) shows that $g(y_{k-1}+\Delta_k) \le g(y_{k-1}) - K_{\mathrm{prog}}\epsilon^{3/2}\rho^{-1/2}$ at each successful iteration, so the last iteration of Algorithm 1 is the only unsuccessful one.

Let k^* be the index of the final iteration with model f_{k^*}, $s_{k^*} = \operatorname{argmin} f_{k^*}$, $A_{k^*} = \nabla^2 g(y_{k^*-1})$, and let $b_{k^*} = \nabla g(y_{k^*-1})$. Since the final iteration is unsuccessful, Lemma 6.1 implies $\|s_{k^*}\| \le \sqrt{\epsilon/(9\rho)}$. Let Δ_{k^*} be the point produced by the call to Solve-final-subproblem, and let $y_{\mathrm{out}} = y_{k^*} + \Delta_{k^*}$ denote the output of Algorithm 1. Note that Solve-final-subproblem guarantees that

$\|\nabla f_{k^*}(\Delta_{k^*})\| \le \epsilon/2$. Moreover, by the same argument we use in the proof of Lemma 6.1, η satisfies Assumption A. Since Assumption B is also satisfied, we have by Corollary 2.5 that $\|\Delta_{k^*}\| \le \|s_{k^*}\|$. Therefore, by Assumption C we have that

$$\nabla^2 g(y_{\text{out}}) \succeq A_{k^*} - 2\rho\|\Delta_{k^*}\| I \succeq -\sqrt{\rho\epsilon} I,$$

where we used $A_{k^*} \succeq -\rho\|s_{k^*}\| I$ and $\rho\|\Delta_{k^*}\| \le \rho\|s_{k^*}\| \le \sqrt{\rho\epsilon}/3$. That is, the output y_{out} satisfies the second condition (24).

It remains to show that $\nabla g(y_{\text{out}})$ is small. Using $\nabla f_{k^*}(\Delta_{k^*}) = b_{k^*} + A\Delta_{k^*} + \rho\|\Delta_{k^*}\|\Delta_{k^*}$ we have

$$\|b_{k^*} + A\Delta_{k^*}\| \le \|\nabla f_{k^*}(\Delta_{k^*})\| + \rho\|\Delta_{k^*}\|^2.$$

Recalling that $\nabla^2 g$ is 2ρ-Lipschitz continuous (Assumption C) we have [18, Lemma 1]

$$\|\nabla g(y_{\text{out}}) - (b_{k^*} + A\Delta_{k^*})\| \le \frac{2\rho}{2}\|\Delta_{k^*}\|^2 \ .$$

Combining the last two displays and using $\|\nabla f_{k^*}(\Delta_{k^*})\| \le \epsilon/2$ and $\rho\|\Delta_{k^*}\|^2 \le \rho\|s_{k^*}\|^2 \le \epsilon/9$ yields

$$\|\nabla g(y_{\text{out}})\| \le \|\nabla f_{k^*}(\Delta_{k^*})\| + 2\rho\|\Delta_{k^*}\|^2 \le \epsilon,$$

which completes the proof of ϵ-second-order stationarity (24) of y_{out}.

We now bound the total number of gradient descent iterations Algorithm 1 uses. Noting that $d/\delta' = K_{\max} d/\delta > 1$ and that $1/(\eta\rho\|\hat{s}\|) > \beta/\sqrt{\rho\epsilon} \ge 1$, we see that a call to Solve-subproblem performs at most $O(1)\beta\rho^{-1/2}\epsilon^{-1/2}(\log(d/\delta) + \log K_{\max} + \log(\beta/\sqrt{\rho\epsilon}))$ iterations. From $\epsilon \le \rho^{1/3}$ we have that $K_{\max} \le 324\,\lceil g(y_0) - g^\star \rceil\,\sqrt{\rho}\epsilon^{-3/2}$, and so the number of iterations has the further upper bound

$$O(1) \cdot \frac{\beta}{\sqrt{\rho\epsilon}} \log\left(\frac{d}{\delta} \cdot \frac{\beta\,\lceil g(y_0) - g^\star \rceil}{\epsilon^2}\right).$$

Multiplying this bound by the upper bound on $K_{\max}$ shows that the total number of gradient steps made by all calls Solve-subproblem is bounded by (25).

Finally, standard analysis [17, Ex. 1.2.3] of gradient descent on smooth funcionts shows that Solve-final-subproblem, which we call exactly once, terminates after at most $2\frac{f(x_0)-f(s_{k^*})}{\eta(\epsilon/2)^2} \le 80\beta(g(y_0) - g^\star)\epsilon^{-2}$ iterations, as f_{k^*} is $\beta + 2\rho R$-smooth and $\eta \le \frac{1}{\beta+2\rho R}$.

References

[1] N. Agarwal, Z. Allen-Zhu, B. Bullins, E. Hazan, and T. Ma. Finding approximate local minima for nonconvex optimization in linear time. *arXiv:1611.01146 [math.OC]*, 2016.

[2] T. Bianconcini, G. Liuzzi, B. Morini, and M. Sciandrone. On the use of iterative methods in cubic regularization for unconstrained optimization. *Computational Optimization and Applications*, 60(1):35–57, 2015.

[3] Y. Carmon, J. C. Duchi, O. Hinder, and A. Sidford. Accelerated methods for non-convex optimization. *arXiv:1611.00756 [math.OC]*, 2016.

[4] C. Cartis, N. I. M. Gould, and P. L. Toint. Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results. *Mathematical Programming, Series A*, 127:245–295, 2011.

[5] C. Cartis, N. I. Gould, and P. L. Toint. Complexity bounds for second-order optimality in unconstrained optimization. *Journal of Complexity*, 28(1):93–108, 2012.

[6] A. R. Conn, N. I. M. Gould, and P. L. Toint. *Trust Region Methods*. MPS-SIAM Series on Optimization. SIAM, 2000.

[7] J. B. Erway and P. E. Gill. A subspace minimization method for the trust-region step. *SIAM Journal on Optimization*, 20(3):1439–1461, 2009.

[8] R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Proceedings of the Twenty Eighth Annual Conference on Computational Learning Theory*, 2015.

[9] N. I. M. Gould, S. Lucidi, M. Roma, and P. L. Toint. Solving the trust-region subproblem using the Lanczos method. *SIAM Journal on Optimization*, 9(2):504–525, 1999.

[10] N. I. M. Gould, D. P. Robinson, and H. S. Thorne. On solving trust-region and other regularised subproblems in optimization. *Mathematical Programming Computation*, 2(1):21–57, 2010.

[11] E. Hazan and T. Koren. A linear-time algorithm for trust region problems. *Mathematical Programming, Series A*, 158(1):363–381, 2016.

[12] N. Ho-Nguyen and F. Kılınç-Karzan. A second-order cone based approach for solving the trust-region subproblem and its variants. *arXiv:1603.03366 [math.OC]*, 2016.

[13] J. Kuczynski and H. Wozniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, 13(4): 1094–1122, 1992.

[14] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. Gradient descent converges to minimizers. In *Proceedings of the Twenty Ninth Annual Conference on Computational Learning Theory*, 2016.

[15] K. Y. Levy. The power of normalization: Faster evasion of saddle points. *arXiv:1611.04831 [cs.LG]*, 2016.

[16] C. Musco and C. Musco. Randomized block Krylov methods for stronger and faster approximate singular value decomposition. In *Advances in Neural Information Processing Systems 28*, 2015.

[17] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Kluwer Academic Publishers, 2004.

[18] Y. Nesterov and B. Polyak. Cubic regularization of Newton method and its global performance. *Mathematical Programming, Series A*, 108:177–205, 2006.

[19] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.

[20] B. A. Pearlmutter. Fast exact multiplication by the Hessian. *Neural computation*, 6(1):147–160, 1994.

[21] N. N. Schraudolph. Fast curvature matrix-vector products for second-order gradient descent. *Neural computation*, 14(7):1723–1738, 2002.

[22] P. D. Tao and L. T. H. An. A D.C. optimization algorithm for solving the trust-region subproblem. *SIAM Journal on Optimization*, 8(2):476–505, 1998.

[23] L. N. Trefethen and D. Bau III. *Numerical Linear Algebra*. SIAM, 1997.